

jc780 U.S. PTO
03/15/00

3-11-00

PTO/SB/05 (2/98) (modified)

Approved for use through 9/30/2000, OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

NEW UTILITY PATENT APPLICATION TRANSMITTAL <small>(only for new nonprovisional applications under 37 CFR 1.53(b))</small>		Attorney Docket Number	4467 US
		First Named Inventor	Peter Newman
		Total Pages in this Submission	61
		Express Mail Label No.	EL263546932US

jc678 U.S. PTO
09/526980

03/15/00

APPLICATION ELEMENTS		ACCOMPANYING APPLICATION PARTS	
<p>1. <input checked="" type="checkbox"/> Fee Transmittal Form (in duplicate)</p> <p><input checked="" type="checkbox"/> Checks Enclosed (\$380.00 / Filing Fee & \$342.00 / Extra Claims)</p> <p>2. <input checked="" type="checkbox"/> Specification (preferred arrangement set forth below)</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Descriptive Title of the Invention <input checked="" type="checkbox"/> Cross Reference(s) to Related Case(s) <input checked="" type="checkbox"/> Statement Regarding Fed sponsored R & D <input checked="" type="checkbox"/> Background of the Invention <input checked="" type="checkbox"/> Brief Summary of the Invention <input checked="" type="checkbox"/> Brief Description of the Drawing(s) <input checked="" type="checkbox"/> Detailed Description <input checked="" type="checkbox"/> Claim or Claims <input checked="" type="checkbox"/> Abstract of the Disclosure <p>3. <input checked="" type="checkbox"/> Drawing(s) (when necessary per 35 USC 113)</p> <p>4. Oath or Declaration <ul style="list-style-type: none"> a. <input checked="" type="checkbox"/> New Declaration <input checked="" type="checkbox"/> Executed b. <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d)) (for continuation/divisional with Box 17 completed) i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b). </p> <p>5. <input type="checkbox"/> Incorporation by Reference (useable if Box 4b is checked). The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.</p>		<p>6. <input type="checkbox"/> Assignment & Assignment Recordation Cover Sheet</p> <p>7. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed)</p> <p>8. <input type="checkbox"/> Information Disclosure Statement & PTO-1449 <ul style="list-style-type: none"> <input type="checkbox"/> Copies of IDS Citation(s) </p> <p>9. <input type="checkbox"/> Preliminary Amendment</p> <p>10. Small Entity Statement <ul style="list-style-type: none"> <input checked="" type="checkbox"/> New Statement enclosed <input type="checkbox"/> Statement filed in prior application. Status still proper and desired </p> <p>11. <input checked="" type="checkbox"/> Return Postcard</p> <p>12. <input type="checkbox"/> _____</p> <p>13. <input type="checkbox"/> _____</p> <p>14. <input type="checkbox"/> _____</p> <p>15. <input type="checkbox"/> _____</p>	
<p>ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, D.C. 20231 </p>			
<p>17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment.</p> <p><input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No: _____ / _____</p> <p>Prior application information: Examiner: _____ Group/Art Unit: _____</p>			

18. CORRESPONDENCE ADDRESS

NAME	Renée M. DuBord, Esq. Fenwick & West LLP			
ADDRESS	Two Palo Alto Square			
CITY	Palo Alto	STATE	CA	ZIP CODE
COUNTRY	U.S.A.	TELEPHONE	(650) 858-7880	FAX (650) 494-1417
Name (Print/Type)	Renée M. DuBord, Esq.		Registration No. (Attorney/Agent)	42,500
Signature	Renée DuBord		Date	3/15/00

EL263546932US

**VERIFIED STATEMENT CLAIMING SMALL ENTITY STATUS
(37 CFR 1.9(f) & 1.27(c))--SMALL BUSINESS CONCERN**

 Docket Number (Optional):
4467 US
Applicant or Patentee: Peter Newman & Pawan GoyalApplication or Patent No.: not yet knownFiling Date or Issue Date: not yet knownTitle: ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES

I hereby declare that I am

 the owner of the small business concern identified below: an official of the small business concern empowered to act on behalf of the concern identified below:NAME OF SMALL BUSINESS CONCERN Ensim CorporationADDRESS OF SMALL BUSINESS CONCERN 1215 Terra Bella Avenue, Mountain View, CA 94043

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.12, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees to the United States Patent and Trademark Office, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention described in:

 the specification filed herewith with title as listed above. the application identified above. the patent identified above.

If the rights held by the above identified small business concern are not exclusive, each individual, concern or organization having rights in the invention must file separate verified statements averring to their status as small entities, and no rights to the invention are held by any person, other than the inventor, who would not qualify as an independent inventor under 37CFR 1.9(c) if that person made the invention, or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d), or a nonprofit organization under 37 CFR 1.9(e).

Each such person, concern or organization having any rights in the invention is listed below:

 No such person, concern, or organization exists. Each such person, concern or organization is listed below:

Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING Srinivasan KeshavTITLE OF PERSON IF OTHER THAN OWNER Chief Technical OfficerADDRESS OF PERSON SIGNING 1215 Terra Bella Avenue, Mountain View, CA 94043SIGNATURE S/KDATE 3/6/00

ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES

INVENTORS

Peter Newman

Pawan Goyal

Prepared by:

Renee M. DuBord

Reg. No. 42,500

Fenwick & West LLP
Two Palo Alto Square
Palo Alto, CA 94306

Express Mail No.: EL263546932US

ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES

INVENTORS

Peter Newman and Pawan Goyal

5 CROSS REFERENCE TO RELATED APPLICATIONS

This application is related to U.S. Patent Serial No. 09/452,286, entitled "Providing Quality of Service Guarantees to Virtual Hosts", by Pawan Goyal et al., filed November 11, 1999, and commonly assigned with the present application. The subject matter of related application U.S. Patent Serial No. 09/452,286 is incorporated by reference herein in its entirety.

10 BACKGROUND

Field of Invention

The present invention relates generally to providing private networking services, e.g. Intranet services, remotely, and more particularly, to allowing a service provider to locate and manage private network servers at the Service Provider's location, while connecting the servers 15 to a customer's premises such that they appear to be local and private to the customer.

Background of the Invention

Networked computer resources are growing more popular as the benefits of sharing computing resources becomes evident. One of the fastest-growing segments of the Internet is the private network market. A private network is an interconnected group of computing resources 20 accessible only by the network members. Security protocols are used to ensure that only

authorized users have access to the network's resources, even where the network operates on the public network infrastructure and protocols. A private network, often belonging to a corporation, may be used to store web pages and other shared information. This information is maintained in a private space, generally screened off from external sources such as the Internet by a firewall 5 which blocks unauthorized access. Such private networks are often referred to as intranets, local area networks (LAN) or wide area networks (WAN).

In a typical LAN implementation, a single location contains a group of individual user computers, as well as one or more dedicated host computers executing server programs to maintain the network's shared information. The private LAN is screened off from the Internet 10 by a firewall, though users may access the Internet if needed. Network traffic intended for the network is allowed through the firewall only if authorized. The resources within the LAN all may communicate using private addresses. It is not necessary to use registered IP addresses for each resource because the system is screened off from the global Internet.

This model may be extended to multiple location sites. Computer networks that span 15 relatively large geographical distances are typically referred to as WANs. In a private WAN, individual sites must be connected in a secure manner. A secure connection between WAN sites may be accomplished using a virtual private network. A virtual private network utilizes ordinary Internet protocols and may also use public communications mediums to connect; however, privacy is ensured through features such as tunneling (data encapsulation) or the use of leased 20 lines. A leased line is a permanent connection between two points that is always active. Resources on a WAN may communicate using only private addresses because the network is screened off from the global Internet.

Private networks contain common elements. Each generally contains a dedicated local server to maintain the shared private network data, and a communications system for providing data communication services between machines on the private network. Communication takes place using a private address space. Because the address spaces for individual private networks need only be locally unique, the address spaces among several different private networks may overlap because the networks are isolated from each other. More specifically, in private intranets, two unrelated intranets at different companies may use the same local addresses for user computers. No conflict arises since the networks are not connected.

Data communications services and servers are not easy to configure, manage, and maintain. Thus, there is an incentive for the service providers that offer access to communications facilities to provide such private network services and servers as well, thereby relieving corporations from the burden of providing these services directly. Some examples of service providers are: Internet Service Providers (ISP), Application Service Providers, Network Service Providers, and Competitive Local Exchange Carriers (CLEC).

It is not economically feasible for a service provider to remotely manage servers located on a customer's premises, and support many different customers in this fashion. Rather a service provider would prefer to offer private network services to multiple customers while keeping all of the server host computers within a location of the service provider for ease of management. Accordingly, service providers typically dedicate a physical host computer as each individual customer's server, and maintain each host computer in the centralized facility. However, this means the service provider will have to own and maintain potentially large numbers of physical host computers, at least one for each customer's server or private network. However, many customers will neither require nor be amenable to paying for the use of an entire host computer.

Generally, only a fraction of the processing power, storage, and other resources of a host computer will be required to meet the needs of an individual customer.

Alternatively, a service provider may utilize one physical host computer to provide commercial host services to multiple customers. Using Transmission Control Protocol (TCP) and other transport protocols, a server application executing on a single physical host can be programmed to process requests made to multiple network addresses. Such functionality is known as virtual hosting.

In virtual hosting, each customer is assigned a network address (or domain name), and is provided with resources on a single, physical host computer, effectively sharing the host with other customers. A client computer requests data from a specific customer's host by targeting communication requests to the appropriate network address (or domain name). The virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer, servicing requests made to a plurality of network addresses and domain names by multiple customers.

However, virtual hosting as it is commonly performed today does not provide many of the beneficial features of private networks. Service providers will have to be able to provide certain features of private networks before customers will be willing to outsource services related to the operation and maintenance of their private network. First, customers will want to ensure that their private data is inaccessible to other customers sharing the same host computers. For instance, if a service provider provides email outsourcing for both Company A and Company B on the same computer, Company A will want to ensure that the directories in which its email is stored are not accessible to Company B, and vice versa.

Additionally, customers will want to ensure that their services are not compromised due to problems originating with another customer. If a service provider uses a single host computer to provide server resources for both Company A and Company B, steps must be taken to ensure that overuse of the resources by Company A does not impact Company B's service.

5 Additionally, faults, crashes, or similar problems caused by one customer must not compromise the service provided to another customer. Such performance degradation issues must be contained by the service provider to impact only the customer responsible for the problem, and not to impact any other customers.

Finally, companies A and B will want their servers to have IP addresses that belong to their own private address spaces. Using addresses from each company's own private address space offers more security because private IP addresses are not reachable over the public Internet. The use of private IP addresses guards against private servers becoming accessible from the public Internet by accidental misconfiguration of equipment. Also, public IP addresses are a limited resource; there is insufficient address space for private networks to consume addresses from the public address space. Furthermore, if a company connects a number of servers or corporate locations together they will wish to establish a virtual private network. To run the routing protocols required for communications within a virtual private network requires a coherent addressing scheme. Such a virtual private network is easier to manage from a private address space.

20 However, the use of private address schemes creates difficulties for the service provider. The service provider may now have several virtual servers assigned to the same IP address, because companies A and B may have overlapping address spaces, as is typical in private

networks. This address overlap can cause the communication network to fail. Typically, one of the virtual servers would become unreachable due to this address overlap error.

Thus in order to satisfy customers' needs, a service provider desiring to provide private network services must be able to guarantee four different kinds of isolation. Functional isolation separates the data and functionality of each customer. Fault isolation protects one customer from the faults created by another customer. Performance isolation allows each customer to receive a performance commitment independent of the behavior of other customers. Address isolation allows each customer to choose the virtual server IP address that it wants to be associated with, independent of other customers.

Virtual hosting currently cannot provide these beneficial features of ordinary private servers. This is due to the inability of a virtual host to allocate appropriate amounts of computer resources of the physical host computer to servicing client requests made to specific virtual hosts, and hence to specific customers. A private virtual server, by contrast, is able to provide the functional, fault, and performance isolation that an ordinary virtual server cannot. A method for creating such a private virtual server is disclosed in the related application identified above, U.S. Patent Serial No. 09/452,286, entitled "Providing Quality of Service Guarantees to Virtual Hosts."

However, a method and system is still needed to allow customers to use their own private address spaces to communicate with a remotely-located private virtual server maintained by a service provider, where the private virtual server addresses may overlap.

SUMMARY OF THE INVENTION

The present invention allows providers of virtual servers to properly differentiate and route transmissions using private addresses on a common host server. The term “private virtual server” as used herein is a virtual server that supports a private address space wherein the private address spaces of different private virtual servers may overlap.

Customers exchange privately-addressed transmissions with a service provider using tunnels to traverse the local or regional network connecting the customer with the service provider. The service provider receives the transmissions at a gateway into the service provider’s data center. The service provider then routes the transmission to the private virtual server belonging to the customer that sent the transmission. The service provider also routes privately-addressed transmissions back to individual customers using tunnels. In this way, the service provider is able to implement a separate routing context on behalf of each individual customer.

The present invention allows for flexibility in designing a private virtual server that suits an individual customer’s needs. For example, an individual customer may have multiple physical sites all utilizing the same private virtual server. Optionally, an individual customer may be assigned more than one private virtual server, for instance, if different divisions of an organization each wish to maintain their own private servers while making all the data available to the organization at large.

In one embodiment, the present invention comprises a multiplexing/demultiplexing mechanism capable of routing signals to and from one or more private virtual servers located on a physical host computer. Each private virtual server is associated with only one private network. The multiplexing/demultiplexing mechanism receives a privately addressed

transmission and routes it to the private virtual server with which it is associated. The privately addressed transmission may then be processed within the private virtual server in the same manner as if the server was actually a physical server resident within a LAN. The multiplexing/demultiplexing mechanism also receives an outgoing transmission from a private virtual server and routes it back to the private network associated with it. Incoming and outgoing transmissions are addressed using a tunneling scheme. Tunneling allows privately-addressed transmissions to be transported over a network that uses global addresses and allows customers to address their private virtual server using a non-unique private address. As a result, all transmissions to and from a customer are sent only to the customer's respective private virtual server.

In another embodiment, the present invention includes more than one physical host computer. Each physical host computer contains one or more private virtual servers. Privately-addressed incoming transmissions are sent to the service provider's data center. The transmissions are placed on tunnels to traverse the local or regional network connecting the customer with the service provider. A tunnel switching mechanism is used to forward transmissions received at the service provider's data center to the proper physical host computer.

In another embodiment, the present invention is a method for locating and managing private network services in a data center location remote from private network users. The method comprises receiving a transmission addressed using a private address of a recipient, and routing the transmission to a private virtual server. The correct private virtual server is determined from the address of either the sender or the recipient of the transmission.

The foregoing merely summarizes aspects of the invention. The present invention is more completely described with respect to the following drawings and detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is an illustration of a private virtual server system in accordance with the present invention.

5 Fig. 2 is an illustration of an embodiment of a private virtual server system with individual tunnels between each customer and associated private virtual server.

Fig. 3 is an illustration of an embodiment of a private virtual server system with a multiplexing/demultiplexing mechanism.

10 Fig. 4 is a block diagram illustrating the tunnel-to-private virtual server traffic flow of a multiplexing/demultiplexing mechanism.

Fig. 5 is a block diagram illustrating the private virtual server-to-tunnel traffic flow of a multiplexing/demultiplexing mechanism.

Fig. 6 is an illustration of an embodiment of a private virtual server system with multiplexing/demultiplexing mechanisms and a tunnel switch.

15 Fig. 7 is a block diagram illustrating the traffic flow of a tunnel switch.

Fig. 8 is a diagram of a customer lookup table for a tunnel switch in an embodiment of the present invention.

Fig. 9 is a diagram of a set of customer forwarding tables for a tunnel switch in an embodiment of the present invention.

20 Fig. 10A is an illustration of an embodiment of a system for the separation of Internet and private virtual server traffic.

Fig. 10B is an illustration of another embodiment of a system for the separation of Internet and private virtual server traffic.

Fig. 11 is an illustration of an embodiment of a private virtual server system using Asynchronous Transfer Mode (ATM) layer two tunneling.

5 Fig. 12 is a block diagram illustrating the traffic flow of an ATM cell from a tunnel to a
private virtual server.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Reference will now be made in detail to several embodiments of the present invention, examples of which are illustrated in the accompanying drawings. Wherever practicable, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

5 The figures depict preferred embodiments of the present invention for purposes of illustration only. One skilled in the art will readily recognize from the following discussion that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles of the invention described herein.

Fig. 1 is an illustration of an embodiment of a private virtual server system, which replaces multiple private Intranet servers. A private virtual server system 100 includes a service provider gateway 152 that connects to a service provider data center 150. System 100 is connected to multiple customer sites 110 via an access network 120. A set of customers, e.g. 112 and 114, are located in potentially any global location, although frequently the set of customers are located within the regional access network of the service provider. As will be evident to one of skill in the art, any number of customers may be supported by the private virtual server system 100. Only the resources of the physical servers owned by the service provider limit the number of customers supported. Furthermore, a service provider may add additional physical servers as required to support additional customers.

As used herein, the terms “customer”, “user”, and “private network user” refer to 20 individuals or groups of individuals accessing the same private virtual server. Typically, a private virtual server “user” is a group of individuals with a shared association. For example, “user” may collectively refer to the employees of a company, or to certain employees within a

division of a company. One company (a “customer”) may have several different users, each corresponding to a different group within the company. Additionally, a “user” may also refer to a single individual.

In the present invention, customers send and receive data using a network. In one embodiment, the protocol used within each customer’s network will be IP (Internet Protocol). The IP addressing format will also be assumed herein for purposes of illustration. However, it will be evident to one of skill in the art that a different network protocol could be used instead of IP within a customer’s private network, for example, Open Systems Interconnection (OSI), Internetwork Packet Exchange (IPX), System Network Architecture (SNA), and Asynchronous Transfer Mode (ATM).

Data from customers 112 and 114 may optionally be aggregated 130 before being sent out over a local or regional network 140. Aggregation 130 is not required, but generally is done to concentrate data traffic for more efficient transmission over a network. For example, a customer site may aggregate the multiple data streams coming from various different users (clients) within the site before transmitting data onto a public network.

Aggregation 130 may also take place between customers and non-customers. This may occur when data transmission sites are located physically close together, to consolidate traffic flow onto the local or regional network 140. In another example, aggregation is also common in digital subscriber line (DSL) systems. A digital subscriber line access multiplexer (DSLAM) concentrates data traffic from multiple DSL loops onto the backbone network before connection to the rest of the local or regional network 140. Data directed to different destinations will be separated and individually routed within the local or regional network 140.

The local or regional network 140 connects a customer with a private virtual server system 100. The local or regional network 140 may be comprised of different types of interconnected networks. System 100 is capable of functioning with a wide variety of different local or regional networks 140, as will be evident to one of skill in the art.

5 Traffic from each individual customer 112 and 114 is aggregated 130 and transported across the local or regional network 140 until it reaches a service provider gateway 152. The service provider gateway 152 provides the connection into the service provider data center 150. The service provider gateway 152 directs traffic from each customer to the private virtual servers they own, and prevents traffic from reaching private virtual servers a customer does not own.

10 The service provider gateway 152 connects to a physical server machine 160. The physical server machine 160 may be any kind of computer adapted to support private virtual servers. It is to be understood that a service provider data center 150 will typically contain more than one physical server machine 160.

15 Located on the physical server machine 160 is a group of private virtual servers 162. In one embodiment, each private virtual server 162 is capable of implementing quality of service guarantees to individual customers. A service provider may implement different quality of service guarantees by allocating different percentages of the physical server machine 160 to servicing each of the private virtual servers 162. Thus the private virtual servers 162 may each consume a different percentage of the resources of physical server machine 160. This resource 20 allocation may be dynamically changed by the service provider as required.

In one embodiment, the IP network address of each private virtual server 162 corresponds to a private address from the private address space of the respective customer assigned to the private virtual server. A private address is any locally-assigned address, which does not have to

be unique within the global Internet. Conversely, a global address is a registered IP address that is unique within the public network system of communications. Generally, connecting a private network to a public network requires network address translation, which maps global addresses onto private addresses at the public/private network boundary. Private addresses are not routable
5 in the public network, therefore transmissions cannot be sent to the private virtual servers 162 across the public network 120 using only their private addresses.

The network addresses for the private virtual servers are chosen to correspond to each customer's private address space. The private address spaces of customers 112 and 114 may overlap. Therefore, the network addresses for these customers' assigned virtual private servers
10 may not be unique. It is possible for private virtual servers 162 to all be assigned the same network address.

Thus the use of private addresses for the private virtual servers creates two different problems: such private addresses cannot be used to transmit data directly across a public network, and there may be address overlap among the private virtual servers.

15 A solution to the problem of transporting privately-addressed transmissions is shown in Fig. 2. Fig. 2 is an illustration of an embodiment of a private virtual server system 200 that uses individual tunnels to connect between each customer and a service provider. The tunnels provide a means to transmit privately-addressed data across a public network.

Fig. 2 shows a private virtual server system 200. The private virtual server system 200
20 includes a service provider gateway 252 that connects to a service provider data center 250. A physical server machine 260 contains three private virtual servers 262. These private virtual servers may have overlapping private IP addresses. Private virtual server 262A is assigned to a customer 220, while private virtual server 262B is assigned to a customer 210. Customers 210

and 220 are connected to the private virtual server system 200 via a local or regional network 240.

Customers 210 and 220 use tunnels to allow privately-addressed data transmissions to traverse the local or regional network 240. Customer 210 has a single tunnel 212 from customer 210's site to the service provider gateway 252. Similarly, customer 220 has a single tunnel 222 from customer 220's site to the service provider gateway 252. A tunnel is created by encapsulating a data transmission within a second type of addressing protocol. Transmissions, which use addresses that are unique only within a narrow scope, may thus be transported across a second network within a wider address scope. In this way, private IP addresses that are only unique within a private IP network are encapsulated for transport across the global public Internet. The origination of the tunnel, on the customer's side, must occur before traffic from multiple customers is aggregated. The numbers of customers and private virtual servers shown in Fig. 2 is merely representative. It is to be understood that a service provider may have more or fewer customers. Furthermore, a physical host machine may contain more or fewer private virtual servers.

System 200 demonstrates how a customer can be communicatively coupled to a private virtual server system using privately-addressed transmissions. However, once privately-addressed transmissions are routed across a public network, the problem of potential address overlap among the destination private virtual servers still remains. Fig. 3 demonstrates a method 20 for routing transmissions to private virtual servers, where the private virtual servers may have overlapping address spaces.

Figure 3 is an illustration of an embodiment of a private virtual server system 300 including a multiplexing/demultiplexing mechanism 350. Private virtual server system 300

provides a separate routing context on behalf of each user to route privately-addressed transmissions between the users and the private virtual servers. System 300 is connected to three customer sites 310, 320 and 330. Each customer uses a tunnel to traverse the local or regional network 340 and arrive at a physical server machine 360 within the private virtual server system 5 300. Customer site 310 uses a tunnel 312, customer site 320 uses a tunnel 322, and customer site 330 uses a tunnel 332. Each tunnel is a different data encapsulation. Multiple tunnels may be carried on the same physical medium connecting to the multiplexing/demultiplexing mechanism 350.

In another embodiment, the tunnels 312, 322, and 332 could be replaced with dedicated 10 leased lines. A leased line is a permanent connection between two points set up by a telecommunications common carrier. A leased line is not part of the global public network, and therefore may carry privately-addressed traffic.

The physical server machine 360 contains three private virtual servers 362, as well as a multiplexing/demultiplexing mechanism 350. Incoming tunnels 312, 322, and 332 15 communicatively couple to the multiplexing/demultiplexing mechanism 350. The private virtual servers 362 may have overlapping IP addresses. A set of internal pointers 314 is used to direct traffic from the multiplexing/demultiplexing mechanism 350 to the correct private virtual server 362, thereby communicatively coupling the multiplexing/demultiplexing mechanism 350 to the private virtual servers 362.

20 The multiplexing/demultiplexing mechanism 350 performs the functions of separating incoming communication streams back into their original constituent streams, and merging multiple separate communication streams onto a single physical communications medium. In one embodiment, the multiplexing/demultiplexing mechanism 350 is implemented in the

network interface card of the physical server machine 360. The multiplexing/demultiplexing mechanism 350 may be implemented in an application specific integrated circuit (ASIC) on the network interface card, or on the software driver.

Incoming tunneled transmissions are sent on a physical medium as packet flows. When 5 the multiplexing/demultiplexing mechanism 350 demultiplexes an incoming set of packet flows, incoming packets are stored in a buffer and one of the fields in the packet header is used to select the incoming packet queue to which the buffer should be linked. An incoming packet queue is a list of pointers wherein each pointer points to a packet buffer. Thus packets arriving with 10 different tunnel identifiers are linked to different incoming packet queues. There is one incoming packet queue per tunnel.

Outgoing transmissions sent by the various private virtual servers are also packet flows. Each private virtual server has its own outgoing packet queue. These packet queues are again implemented as lists of pointers to packet buffers. The scheduler in the network interface card selects which of the non-empty outgoing packet queues it should serve next. The scheduler takes 15 the pointer from the head of the selected queue and transmits the packet from the buffer that is pointed to by the pointer.

The tunnel switching operations of the mutiplexing/demultiplexing mechanism 350 are shown in more detail in Figs. 4 and 5. Fig. 4 is a block diagram illustrating the tunnel-to-private virtual server traffic switching functions of a multiplexing/demultiplexing mechanism. Fig. 5 is 20 a block diagram illustrating the private virtual server-to-tunnel traffic switching functions of a multiplexing/demultiplexing mechanism.

Fig. 4 shows an embodiment of a physical host computer 400 including three private virtual servers 420A, 420B, and 420C. Each private virtual server includes an associated IP

stack 422. An IP stack is a set of software processes that together manage the transfer of information in packets according to Internet protocols. Internet protocols define the rules and conventions for exchanging information across the Internet. The number of private virtual servers shown is merely illustrative. It is to be understood that a physical host computer may 5 include more or fewer private virtual servers.

The physical host computer 400 also includes a multiplexing/demultiplexing mechanism 410. Multiplexing/demultiplexing mechanism 410 has an internal pointer 430 directed at each IP stack 422. Multiplexing/demultiplexing mechanism 410 includes a lookup table 412 that includes a list of incoming tunnel identifiers and pointers to their associated IP stacks within the 10 physical host computer 400. Lookup table 412 is created by software residing on the physical host computer 400, which associates each IP stack with a particular customer, and ensures that transmissions sent from a particular customer are only directed to that particular customer's private virtual server.

A data packet 452 enters the multiplexing/demultiplexing mechanism 410 using an 15 incoming tunnel 450. The tunnel 450 is not an individual physical connection, but is a means of encapsulating the packet 452 to permit routing across a public network. Although only one incoming tunnel 450 is shown in Fig. 4, this is merely illustrative. It is to be understood that there may be multiple different tunnels entering the multiplexing/demultiplexing mechanism 410.

20 The tunnel identifying information of the incoming packet 452 identifies which customer sent the packet 452. Packet 452's incoming tunnel identifying information is stripped 446, and the tunnel identifying information is presented 444 to the lookup table 412. Lookup table 412

returns 442 an internal pointer 430 for the appropriate IP stack 422. Packet 452 is then routed 440 to the identified IP stack 422 using the internal pointer 430.

An alternative embodiment is to embed the internal pointer 430 in the control data structure that controls the operation of the tunnel. The control data structure is typically located 5 on the network interface card. The network interface card is typically located on the physical server machine 400.

Fig. 5 is a block diagram illustrating an embodiment of the private virtual server-to-tunnel traffic switching operations for the physical host computer 400. Multiplexing/demultiplexing mechanism 410 contains a lookup table 512 that consists of a list of 10 IP stacks in the physical host computer 400 and their associated outgoing tunnel identifiers. Multiplexing/demultiplexing mechanism 410 ensures that packets sent from a customer's private virtual server are routed back to that customer on the correct outgoing tunnel.

A packet 552 is originally sent from the IP stack 422 of one of the private virtual servers 420. The packet 552 is directed 530 to the multiplexing/demultiplexing mechanism 410. Upon 15 arrival at the multiplexing/demultiplexing mechanism 410, the internal tunnel IP stack identifier is read 540, and this identifier is presented 542 to lookup table 512.

Lookup table 512 returns 544 the associated outgoing tunnel identifying information. The outgoing tunnel identifier is added 546 to packet 552. Packet 552 is then sent out using the appropriate outgoing tunnel 550. The outgoing tunnel 550 is associated with the customer 20 assigned to the private virtual server 420, which originally sent the packet 552. Although only one outgoing tunnel is shown in Fig. 5, this is merely illustrative. It should be understood that there may be multiple different outgoing tunnels. Each tunnel is not a separate physical

connection; it is a specific encapsulation of data allowing the data to be separated out from other data sent on a physical connection.

An alternative embodiment is for the IP stack 422 of each private virtual server 420 to encapsulate the packet 552 with the proper outgoing tunnel header and tunnel identifier, thereby 5 creating the tunnel 550. In this embodiment, the lookup table 512 is not used. The multiplexer/demultiplexer mechanism 410 merges the encapsulated and identified packet streams together, and sends them out.

A service provider will typically own and manage multiple physical server machines. Additionally, customers may wish to purchase more than one private virtual server, for example, 10 for separate divisions within the same company. A customer may also have multiple physical customer sites using the same private virtual server, wherein each customer site uses a different tunnel to communicate with the service provider. A tunnel switch supports these different configurations in a private virtual server system.

A tunnel switch comprises one or more physical interfaces, with each interface capable of 15 carrying many multiplexed tunnels. A tunnel switch will typically be capable of supporting a number of different physical interface technologies and a number of different types of tunnels for each type of physical interface. A tunnel switch performs two separate services: switching tunnels, and switching packets within a tunnel.

In a switching tunnels service, all packets arriving on one incoming tunnel are forwarded 20 to an outgoing tunnel. This may be implemented without looking into the headers of the packets themselves. The tunnel switch notes which tunnel and physical interface the packet arrived on. Given the identifier of the incoming tunnel and the incoming physical interface, the tunnel switch uses a lookup table to specify the outgoing tunnel and physical interface.

In a switching packets service, incoming tunnels are terminated at the tunnel switch. The packets within each tunnel are extracted and switched individually based upon the incoming physical interface, the incoming tunnel, and information from the header of the packet. A private virtual server system supporting multiple physical servers, multiple private virtual servers per 5 customer, and multiple customer sites per private virtual server will support such a switching packets service in addition to supporting a switching tunnels service.

Fig. 6 is an illustration of an embodiment of a private virtual server system containing a tunnel switch. Private virtual server system 600 provides a separate routing context on behalf of each customer to route between the customer and each customer's associated private virtual 10 server. A private virtual server system 600 is communicatively coupled to three customer sites 612, 614, and 616 across a local or regional network 620. The number of customers, physical servers, and private virtual servers shown in Fig. 6 is merely representative. It is to be understood that a service provider may have more or fewer customers and more or fewer 15 physical servers. Furthermore, a physical server may contain more or fewer private virtual servers.

Each customer site is connected to the system 600 via an external tunnel that traverses the local or regional network 620 and communicatively couples to a tunnel switch 632. Customer site 612 uses external tunnel 622, customer site 614 uses external tunnel 624, and customer site 616 uses external tunnel 626. A tunnel switch 632 operates as the gateway to physical servers 20 640 and 650, and supports tunnels on multiple physical interfaces. Tunnel switch 632 is communicatively coupled to physical server 640 via a set of tunnels 634. Tunnel switch 632 is communicatively coupled to physical server 650 via a set of tunnels 636.

Physical server 640 contains a multiplexing/demultiplexing mechanism 642, which is linked via a set of internal pointers 662 to a set of private virtual servers 660. Physical server 650 contains a multiplexing/demultiplexing mechanism 652, which is linked via a set of internal pointers 682 to a set of private virtual servers 680.

5 Fig. 7 is a block diagram of an embodiment of the functions of a tunnel switch. A tunnel switch 700 is connected to a set of physical interfaces 712 and 714. Each physical interface 712 and 714 is capable of carrying a set of external tunnels. Physical interface 712 carries tunnels 710A, 710B, 710C, 710D, 710E and 710F. Physical interface 714 carries tunnels 716A and 716B.

10 Tunnel switch 700 also includes a set of outgoing physical interfaces 732 carrying a set of outgoing tunnels 730 to a set of physical host computers 720. Physical interface 732A carries tunnels 730A and 730B to physical host computer 720A. Physical interface 732B carries tunnel 730C to physical host computer 720B. Physical interface 732C carries tunnels 730D and 730E to physical host computer 720C. The tunnel switch 700 includes a customer lookup table 800 and a set of customer forwarding tables 900. The features of the customer lookup table and 15 customer forwarding tables will be discussed in more detail before fully explaining the functions of the tunnel switch 700.

20 Fig. 8 is an embodiment of a customer lookup table 800, and Fig. 9 is an embodiment of a set of customer forwarding tables 900. Together, table 800 and set of tables 900 are suitable for switching a set of transmissions from an incoming physical interface and tunnel to an outgoing physical interface and tunnel. An incoming transmission may be arriving at the tunnel switch either from a customer, or from a private virtual server. Similarly, an outgoing transmission may be directed towards a private virtual server or a customer. Tables 800 and 900

operate to switch transmissions in both directions. Customer lookup table 800 is used as an index into the correct customer forwarding table from the set of customer forwarding tables 900.

The customer lookup table 800 associates external tunnel and physical interface identifiers with a particular customer. Each customer listed in the customer lookup table 800 has an associated customer forwarding table. For example, customer 1's information is contained in customer forwarding table 910, and customer 2's information is contained in customer forwarding table 920. When a customer communicates with a private virtual server, each data transmission will arrive on a particular physical interface using a particular tunnel.

Customer lookup table 800 contains four fields: incoming physical interface, incoming tunnel identifier, service, and customer identifier. Each incoming physical interface and incoming tunnel identifier entry will reference a unique customer identifier. This customer identifier provides an index to the correct customer forwarding table associated with this physical interface/tunnel identifier pair. For example, a transmission arriving on physical interface 712 and incoming tunnel 710B would be indexed to customer 1's customer forwarding table (table 910 from Fig. 9). In another example in the opposite direction, a transmission arriving on physical interface 732A and incoming tunnel 730B would also be indexed to customer 1's customer forwarding table.

Each customer forwarding table 910, 920 and 930 contains three fields: destination IP address, outgoing tunnel identifier, and outgoing physical interface. Based upon the destination IP address of a particular transmission, the proper outgoing tunnel and outgoing physical interface is determined. Using customer lookup table 910 as an example, transmissions with a destination IP address of the "main server" for customer 1 would be placed on outgoing tunnel 730A on physical interface 732A. In the opposite direction, transmissions with a destination IP

address of “site 1” for customer 1 would be placed on outgoing tunnel 710E on physical interface 712.

The information in the customer lookup tables 900 is segregated by customer because the private address spaces of different customers may overlap, and therefore the destination IP addresses on each individual customer forwarding table are not unique within the set of all customer forwarding tables. For example, “main server” of customer 1 and “server” of customer 2 may have the same IP address.

The service field in the customer forwarding table 800 identifies whether only a tunnel switching service (TS) is required, or a packet switching service (PS) is required. If only a tunnel switching service is required, the associated customer forwarding table will contain only a single entry specifying the proper outgoing tunnel and outgoing physical interface. For example, customer lookup table 800 indexes transmissions arriving on physical interface 712 and incoming tunnel 710C to customer 2, and identifies these transmissions as requiring only a tunnel switching service. Customer lookup table 920 associated with customer 2 directs all traffic to outgoing tunnel 730D on physical interface 732C.

The customer lookup table 800 and customer forwarding tables 900 allow the tunnel switch 700 to support a variety of different private network configurations. For example, customer lookup table 800 shows that customer 1 uses two different incoming tunnels: 710A and 710B. Customer forwarding table 910 shows that customer 1 also has two different destination IP addresses corresponding to two different private virtual servers: “main server” and “backup server.” Additionally, customer 1 has two different destination IP addresses corresponding to two different customer sites: “site 1” and “site 2.” However, customer forwarding table 920 shows that customer 2 has only one private virtual server (“server”). The

example private network configurations referred to herein are merely illustrative. It will be understood by one of skill in the art that many different configurations are possible.

Referring back to Fig. 7, Fig. 7 shows the steps associated with switching a packet between an external tunnel (in one embodiment, from a customer site) and a private virtual server. A physical interface 712 connects a set of external tunnels 710 to the tunnel switch 700. A packet 718 arrives on one of the external tunnels 710B. The incoming physical interface and tunnel identification information is read 740 from the packet 718, and presented 744 to a customer lookup table 800.

The customer lookup table 800 uses the physical interface and tunnel identifier to return 748 the correct customer forwarding table (910) for use with the packet 718. A group of customer forwarding tables 900 contains a customer forwarding table 910 that is associated with the customer that uses physical interface 712 and tunnel 710B. Packet 718's destination IP address, for example, "main server" is presented 750 to customer forwarding table 910.

From the information contained in customer forwarding table 910, the correct outgoing physical interface and tunnel identifier for packet 718 is identified 754. Referring to Fig. 9, destination IP address "main server" corresponds to outgoing tunnel 730A and physical interface 732A. Packet 718 is then placed on tunnel 730A that will transport it to the physical host computer 720A.

Customers utilizing a private virtual server system may also be accessing the global Internet. In this case, both the private virtual server communications traffic and the Internet communications traffic are sent out on a public network. A method for separating the two traffic streams is required.

In one embodiment, shown in Fig. 10A, Internet traffic is separated out at the customer's site. Fig. 10A shows a system 1000 for separating Internet and private virtual server traffic. In system 1000, a traffic separation mechanism 1012 is located at the customer site 1010. A tunnel 1022 is created for sending private virtual server traffic across the local or regional network 1020. Internet traffic 1024 is separated out before reaching the local or regional network 1020 and sent separately to the Internet 1030.

Fig. 10B shows another embodiment of a system for separating Internet and private virtual server traffic. System 1050 separates Internet traffic out after all customer traffic has reached a private virtual server system 1040. A customer site 1010 sends both Internet and private virtual server traffic out on a tunnel 1026 across the local or regional network 1020. All traffic arrives at the private virtual server system 1040 through a tunnel switch 1060.

Within the tunnel switch 1060, traffic addressed to the Internet is segregated out from traffic directed to the customer's private virtual server(s). The Internet traffic is sent out to the Internet 1030 via a public communications channel 1064. Internet responses back to the customer are returned in the same way that they were sent. The tunnel switch 1060 sends private virtual server traffic out on an internal tunnel 1062 to the physical server machine 1070 that holds the customer's private virtual server.

A variety of tunneling protocols may be used to create the tunnels used in the present invention. It is preferable to use a layer 2 tunneling protocol for security purposes. Layer 2 tunnels provide solid circuits that make it difficult to "spoof" a fake source address. Packets using a layer 3 tunneling protocol may be created with a fake IP address, allowing a packet to appear to come from a customer when in fact it was sent by an unauthorized third party. Tunnels using a layer 3 tunneling protocol require encryption so that an intruder cannot decode the

information. Encryption is also required to protect against accidentally routing the traffic to an incorrect destination, because the encryption will prevent mis-directed traffic from being decoded. Preference for tunnel type depends upon the service provider's network architecture.

Typical examples of tunnel protocols include Asynchronous Transfer Mode virtual circuits (ATM VCs), frame relay virtual circuits (FR VCs), the Point-to-Point Protocol (PPP) across the Layer 2 Tunneling Protocol (L2TP), and IP security protocol (IPsec). ATM or frame relay virtual circuits may be delivered to a customer using Digital Subscriber Line (DSL) access. However, it will be evident to one of skill in the art that other tunneling protocols may be used to implement the tunnels of the present invention.

Fig. 11 is an illustration of an embodiment of a private virtual server system using ATM layer 2 tunneling with DSL access. A system 1100 spans a customer site 1110, a local or regional network 1140, and a service provider 1150.

The customer site 1110 contains a group of computers 1112A, 1112B, and 1112C. The computers 1112 are representative of the customer's on-site private network. Such a private network could be comprised of more or fewer individual computers, as well as other types of equipment such as storage devices.

Components within a customer's on-site private network communicate using a private IP addressing scheme. A packet 1114 is a privately-addressed transmission sent using the local IP protocol of the customer site 1110. Packet 1114 is sent from customer 1110 to customer 1110's private virtual server. Packet 1114 has a destination address corresponding to a private virtual server of service provider 1150.

The on-site private network of devices 1112 is connected to a customer premise equipment box (CPE) 1120. Packet 1114 is routed to CPE 1120. The CPE connects the customer site 1110 to a tunnel 1132, established between the customer site 1110 and the service provider site 1150. The tunnel functions as a bi-directional data pipe, and is typically established when a customer subscribes to a service provider providing a tunneling service. A fixed ATM tunnel is referred to as a permanent virtual circuit. The CPE divides the data of packet 1114 into cells (or packets) of a fixed size designated by the ATM protocol, and encapsulates them with a Virtual Channel Identifier (VCI) used to direct the cells through the data pipe. An ATM cell 1118 is created from the division and encapsulation of packet 1114.

Tunnel 1132 terminates at a tunnel switch 1160 of the service provider 1150. Tunnel switch 1160 is connected via a set of bi-directional ATM tunnels 1172 to a set of physical servers 1170. In another embodiment, two sets of unidirectional tunnels are used to provide a two-way connection between customer site 1110 and service provider 1150.

Fig. 12 shows an embodiment of the traffic flow of an ATM cell from a tunnel to a private virtual server. The system of Fig. 12 presents in more detail the functions of the service provider data service center 1150 shown in Fig. 11. In Fig. 12, the service provider data service center 1150 contains the tunnel switch 1160 and a physical host computer 1170. It is to be understood that more than one physical host computer 1170 may be included in system 1150.

A set of incoming external tunnels 1132 terminates at the tunnel switch 1160. Tunnel 1132B carries the incoming ATM cell 1118. Referring back to Fig. 11, ATM cell 1118 is created by the ATM CPE 1120 on the customer site 1110. ATM cell 1118 is created from a privately-addressed IP packet 1114 originally sent by one of the computers 1112.

In Fig. 12, the VCI identifying the incoming tunnel of ATM cell 1118 is read 1212, and presented 1214 to a customer lookup table 1230. Customer lookup table 1230 returns 1216 the correct customer forwarding table (1232B). Table 1232B is accessed from a set of customer forwarding tables 1232.

5 Customer forwarding table 1232B returns 1218 the identity of the outgoing tunnel required to reach the correct private virtual server destination (1250A) for ATM cell 1118. ATM cell 1118 is then placed 1220 on an internal tunnel 1172 to reach the private virtual server 1250A. Internal tunnel 1172 terminates at a multiplexing/demultiplexing mechanism 1290 of the physical host computer 1170, which contains private virtual server 1250A.

10 The VCI of ATM cell 1118 is stripped 1252, and the VCI is presented 1254 to lookup table 1260. The rest of the ATM cell 1118, including the original source/destination addresses and payload, is held until the original IP packet 1114 sent by a customer 1112 can be reassembled. The ATM protocol uses relatively small cells compared to other transmission protocols. Therefore multiple ATM cells may be required to reassemble the original pre-ATM-
15 encapsulation packet 1114.

Once the original IP packet 1114 has been reassembled 1262, the location of the IP stack of the destination private virtual server 1250A is retrieved 1256 from lookup table 1260. Packet 1114 is routed 1280 to the IP stack of private virtual server 1250A.

20 Although the invention has been described in considerable detail with reference to certain embodiments, other embodiments are possible. As will be understood by those of skill in the art, the invention may be embodied in other specific forms without departing from the essential characteristics thereof. For example, different protocols may be used to create tunnels.

Accordingly, the present invention is intended to embrace all such alternatives, modifications and variations as fall within the spirit and scope of the appended claims and equivalents.

We claim:

1 1. A system for locating and managing private network services using private addresses

2 in a location remote from private network users, comprising:

3 a host computer executing a plurality of private virtual servers, each private virtual
4 server associated with a private address space and providing private network
5 services to a plurality of private network users located remotely from the
6 private virtual server; and

7 a multiplexing/demultiplexing mechanism executed by the host computer, and
8 communicatively coupled to a network to receive signals from any of the
9 private networks users, and to route a received signal to the private virtual
10 server associated with the private network user.

1 2. The system of claim 1 wherein the multiplexing/demultiplexing mechanism switches

2 signals between private virtual servers and tunnels associated with private network users.

1 3. The system of claim 1, wherein the multiplexing/demultiplexing mechanism

2 demultiplexes an incoming set of signals into segregated signals, and routes each segregated
3 signal to the private virtual server associated with the private network that transmitted the
4 segregated signal.

1 4. The system of claim 1, wherein the multiplexing/demultiplexing mechanism receives

2 an outgoing set of signals, and routes the signals to an outgoing tunnel associated with the
3 private virtual server that transmitted the signals.

1 5. The system of claim 1, wherein the multiplexing/demultiplexing mechanism is located
2 on the host computer.

1 6. The system of claim 1, wherein the multiplexing/demultiplexing mechanism contains
2 a lookup table, the lookup table storing associations between tunnel identifiers identifying
3 tunnels for private networks and private virtual servers that service the private networks.

1 7. A system for locating and managing private network services in a location remote
2 from private network users, comprising:

3 a plurality of host server computers;

4 a plurality of private virtual servers, each private virtual server adapted to execute on
5 one of the plurality of host server computers wherein each private virtual
6 server is associated with a private network user;

7 a switching mechanism communicatively connected to the plurality of host server
8 computers to receive signals from the private network users and route each
9 signal to the host server computer executing the private virtual server
10 associated with the private network user; and

11 a multiplexing/demultiplexing mechanism communicatively coupled to the switching
12 mechanism for receiving signals from the switching mechanism and routing
13 signals to the correct private virtual server executing on the host server.

1 8. The system of claim 7, wherein the signals are addressed using a private address
2 space.

1 9. The system of claim 7, further including:

2 at least one tunnel for transmitting a signal between a customer's private network and
3 the switching mechanism.

1 10. A system for managing virtual servers using private addresses wherein the private
2 addresses overlap, and wherein the virtual servers are located in a location remote from private
3 network users, comprising:

4 a plurality of host computers;

5 a plurality of virtual servers residing on the plurality of host computers;

6 a tunnel switching mechanism communicatively coupled to the host computers, and
7 communicatively coupled to the network to receive signals from any of the
8 private networks users, to route received signals to each host computer
9 executing the virtual server associated with the private network user; and

10 a multiplexing/demultiplexing mechanism executed by each host computer, and
11 communicatively coupled to a network to receive signals from any of the
12 private networks users, and to route received signals to the virtual server
13 executing on the host computer that is associated with the private network
14 user.

1 11. The system of claim 10, wherein the private address spaces of a first and a second
2 virtual server overlap.

1 12. The system of claim 10, further including:

2 a customer lookup table, the customer lookup table storing associations between
3 physical interfaces and tunnel identifiers identifying tunnels for private
4 networks and a plurality of customer forwarding tables; and
5 a plurality of customer forwarding tables, each customer forwarding table associating
6 network addresses with physical interfaces and tunnel identifiers.

1 13. In a system comprising a host computer containing a plurality of virtual servers that

2 each support a private address space wherein the private addresses of two or more of the virtual
3 servers overlap, a method for locating and managing private network services in a data center
4 location remote from private network users of the virtual servers, the method comprising:

5 receiving a transmission addressed using a private address of a recipient; and
6 routing the transmission to a virtual server associated with at least one of the sender
7 or the recipient of the transmission.

1 14. The method of claim 13, further comprising receiving the transmission via a tunnel

2 from a private network.

1 15. The method of claim 13, further comprising routing the transmission based upon a

2 layer two tunnel identifier.

1 16. The method of claim 13, wherein routing the transmission to a virtual server

2 comprises:

3 terminating an incoming tunnel containing a transmission; and

4 multiplexing the transmission to a virtual server.

1 17. The method of claim 16, including:

2 reading a tunnel identifier contained in the transmission; and

3 selecting a virtual server based upon the tunnel identifier.

1 18. The method of claim 13, wherein routing the transmission to a virtual server

2 comprises:

3 terminating an incoming tunnel containing a transmission;

4 switching the transmission to a tunnel connected to a physical host computer

5 containing a customer's virtual server;

6 terminating the tunnel at the physical host computer; and

7 multiplexing the transmission to the virtual server.

1 19. In a system comprising a host computer containing a plurality of virtual servers
2 which support a private address space wherein a first and a second virtual server private address
3 overlap, a method for performing private network services using private addresses in a location
4 remote from private network users, comprising:
5 storing a customer lookup table, the customer lookup table storing associations
6 between physical interfaces and tunnel identifiers identifying tunnels for
7 private networks and a plurality of customer forwarding tables;
8 storing a plurality of customer forwarding tables, the customer forwarding tables
9 associating network addresses with physical interfaces and tunnel identifiers;
10 receiving a transmission on a physical interface, the transmission containing a tunnel
11 identifier;
12 determining the correct customer forwarding table from the customer lookup table
13 using the physical interface and the tunnel identifier;
14 determining via the customer forwarding table a physical interface and tunnel
15 identifier associated with a network address of the transmission; and
16 sending the transmission to the network address on the determined physical interface
17 using the determined tunnel identifier.

1 20. In a system comprising a host computer containing a plurality of virtual servers
2 which support a private address space wherein a first and a second virtual server private address
3 overlap, a method for a private network to use private network services, wherein the private
4 network services are located remotely from the private network, the method comprising:

5 sending a privately addressed transmission on a tunnel to a virtual server; and
6 receiving a privately addressed transmission back from the virtual server.

1 21. The method of claim 20, wherein the privately addressed transmission does not
2 include a registered IP address.

1 22. The method of claim 20, wherein the tunnel encapsulates the privately-addressed
2 transmissions in a layer two protocol.

1 23. The method of claim 20, further comprising:

2 segregating a first transmission including an Internet address from a second
3 transmission including a virtual server address; and
4 sending the second transmission on a tunnel.

1 24. A method for creating a software architecture suitable for implementing a virtual
2 server system using private addresses wherein the private addresses overlap, the method
3 comprising:

4 implementing a tunneling protocol to tunnel privately-addressed transmissions
5 between a plurality of virtual servers and a plurality of users; and

implementing a separate routing context on behalf of each user to route privately-addressed transmissions between the users and the virtual servers.

1 25. The method of claim 24, wherein the tunneling protocol is Asynchronous Transfer
2 Mode virtual circuits.

1 26. The method of claim 24, wherein the tunneling protocol is frame relay virtual
2 circuits.

1 27. The method of claim 24, wherein the tunneling protocol is the Point-to-Point
2 protocol across the Layer two Tunneling Protocol.

1 28. The method of claim 24, wherein the tunneling protocol is the Internet Protocol
2 security protocol.

1 29. A computer program product for switching signals between private virtual servers
2 and tunnels associated with private network users, the computer program product comprising:

3 program code for demultiplexing an incoming set of signals into segregated signals;
4 and

5 program code for routing each segregated signal to the private virtual server
6 associated with the private network user that transmitted the segregated signal.

1 30. The computer program product of claim 29, further including:
2 program code for receiving an outgoing signal; and

3 program code for routing the signal to an outgoing tunnel associated with the private
4 virtual server that transmitted the signal.

1 31. A computer program product for managing virtual servers using private addresses
2 wherein the private addresses overlap, and wherein the virtual servers are located in a location
3 remote from private network users, the computer program product comprising:

4 program code for creating a plurality of virtual servers residing on a plurality of host
5 computers;

6 program code to receive signals from any of the private networks users, and to route
7 received signals to each host computer executing the virtual server associated
8 with the private network user; and

9 program code to route received signals to the virtual server executing on the host
10 computer that is associated with the private network user.

1 32. The computer program product of claim 31, further including:

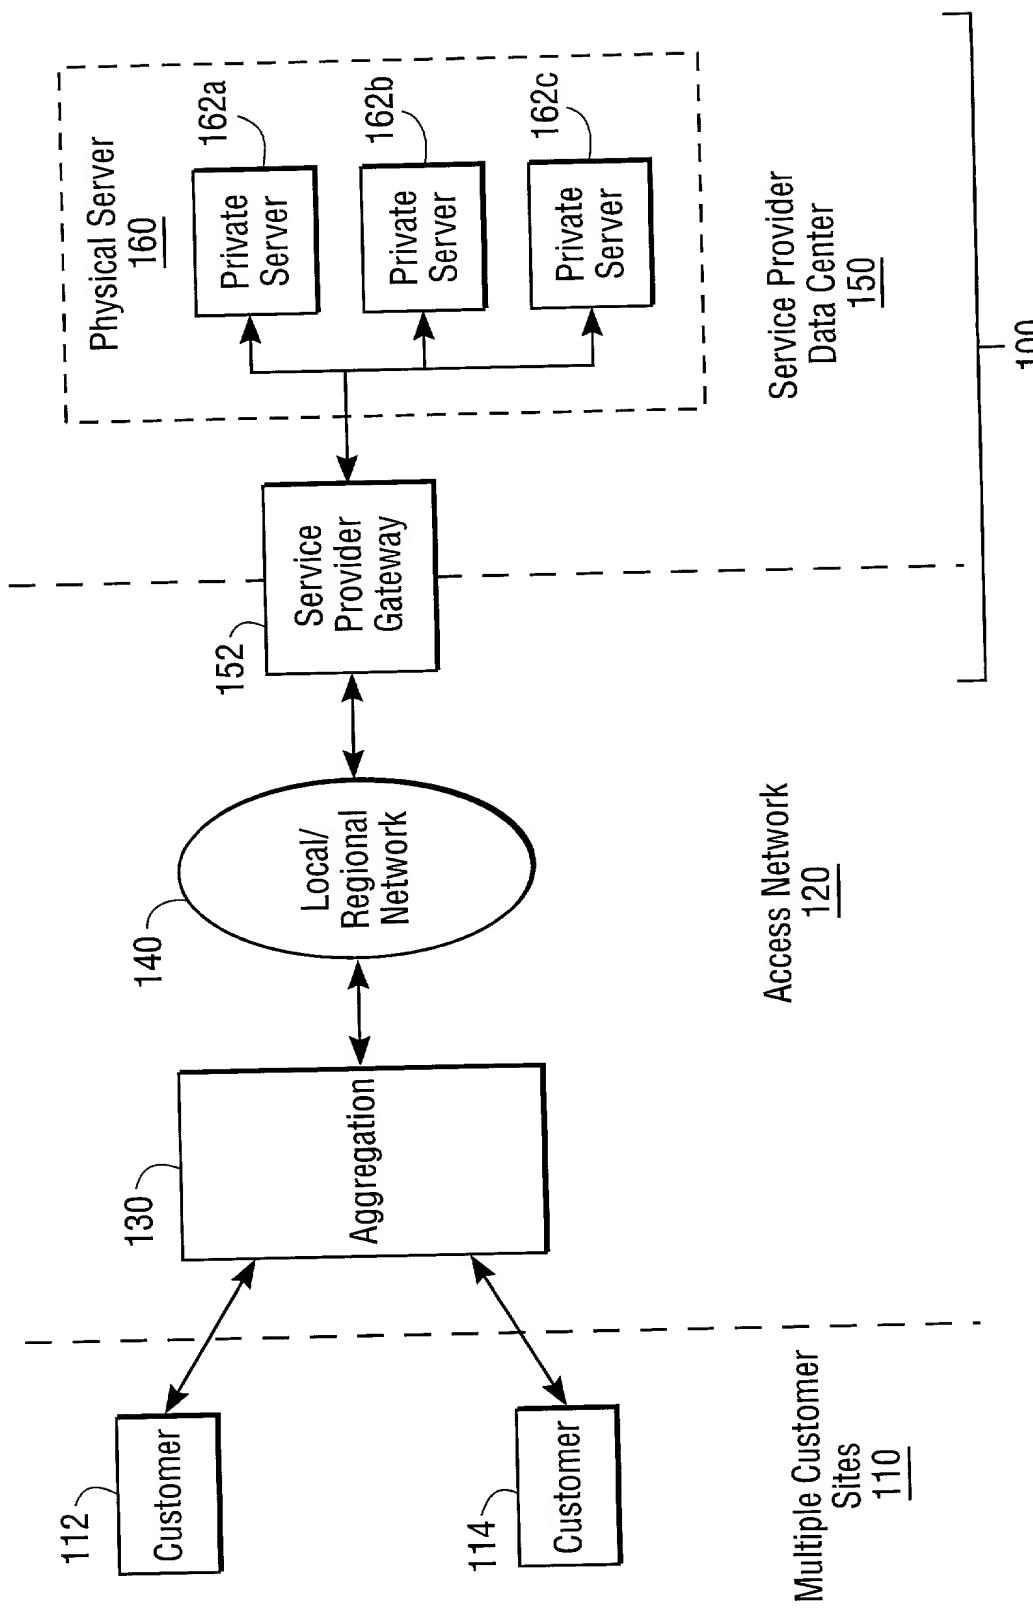
2 program code for storing associations between physical interfaces and tunnel
3 identifiers identifying tunnels for private networks, and a plurality of customer
4 forwarding tables; and

5 program code for creating a plurality of customer forwarding tables, each customer
6 forwarding table associating network addresses with physical interfaces and
7 tunnel identifiers.

ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES

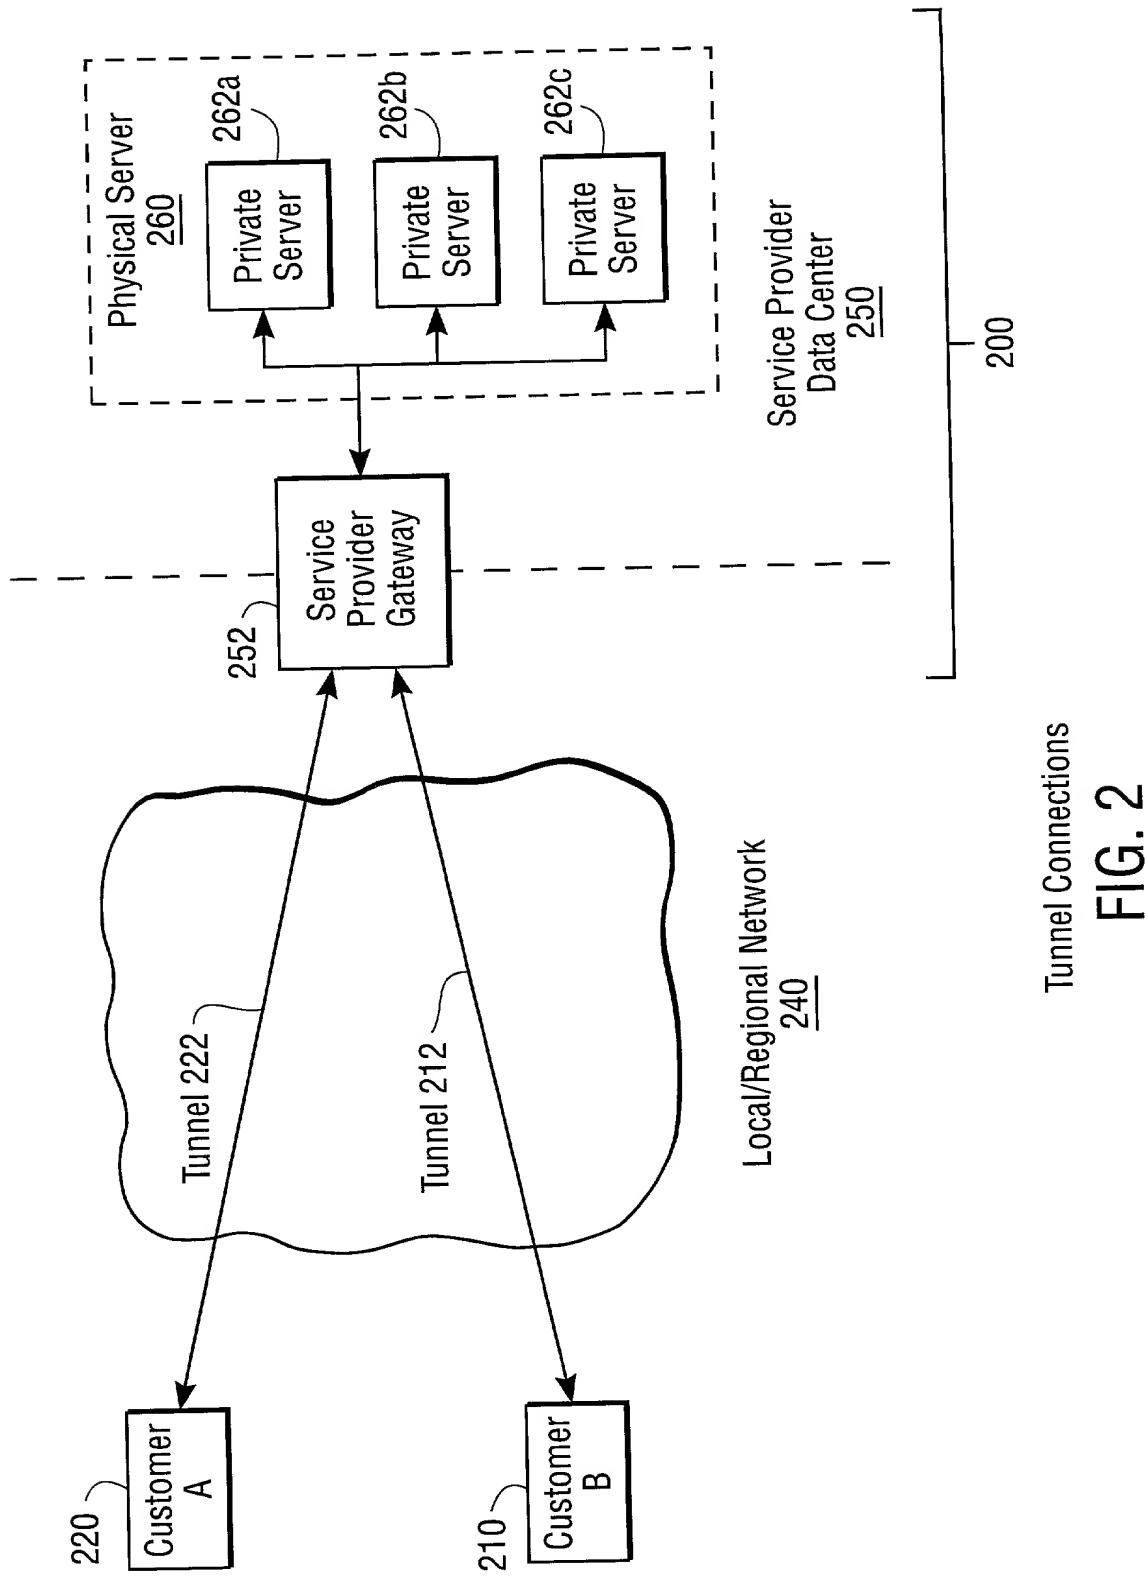
ABSTRACT OF THE DISCLOSURE

A method and system allows a service provider to provide Intranet services remotely by assigning private virtual servers to customers. Each customer addresses transmissions to one or more private virtual servers using private addresses from the customer's private Intranet. The addresses of different private virtual servers do not have to be unique and may overlap. Customers exchange privately-addressed transmissions with the service provider using tunnels to traverse the local or regional network connecting the customer with the service provider. The service provider routes the transmissions to the relevant private virtual server belonging to the customer that sent the transmission. The service provider also routes privately-addressed transmissions back to individual customers using tunnels.



Private Virtual Server Intranet
FIG. 1

220 252 262a 262b 262c 240 200 200



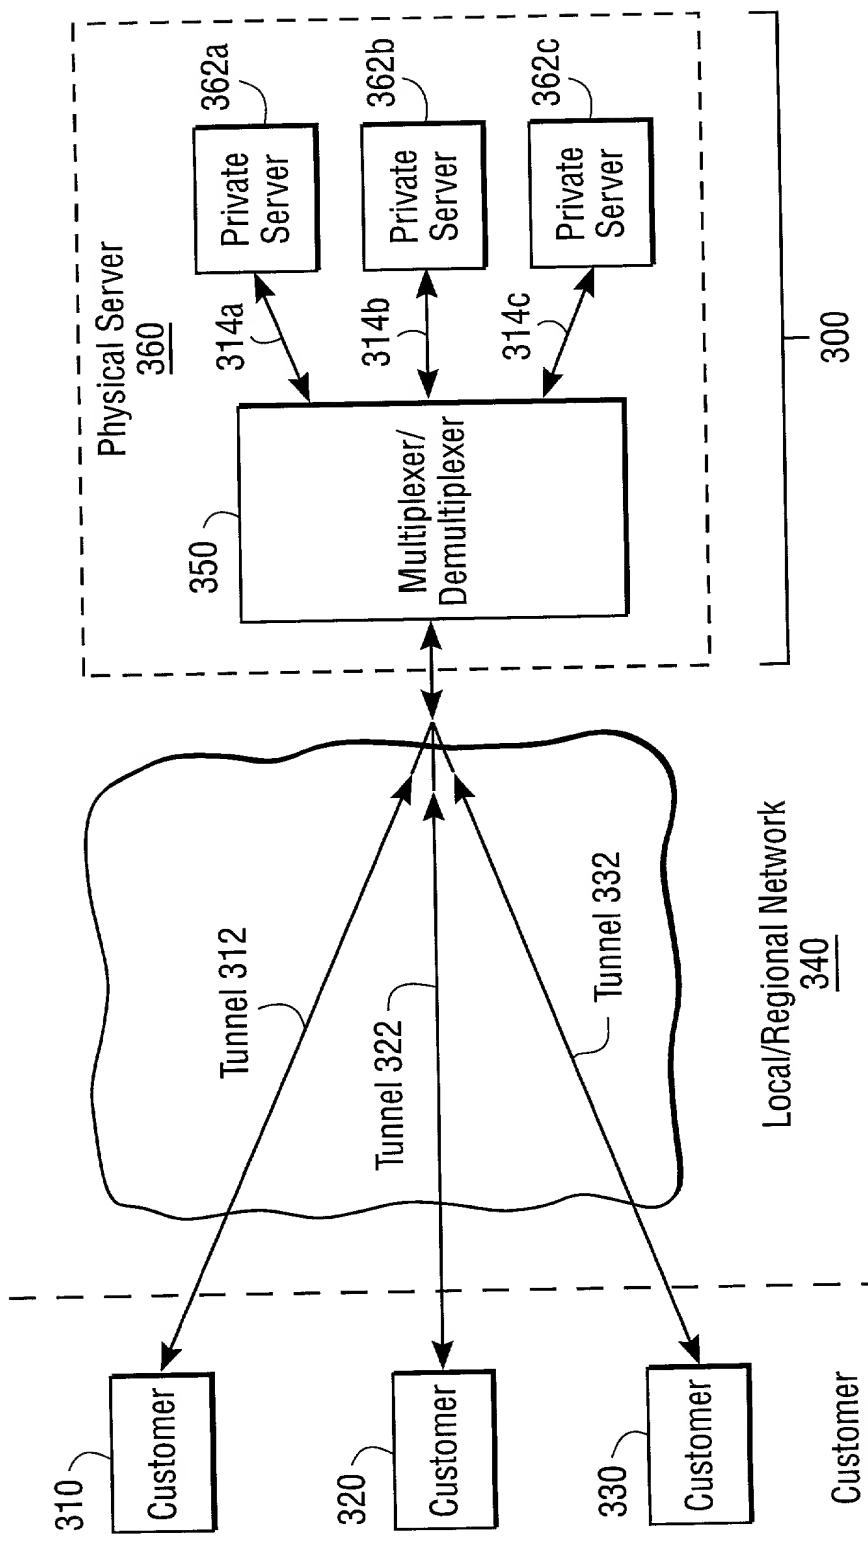


FIG. 3
Multiplexing of Tunnels

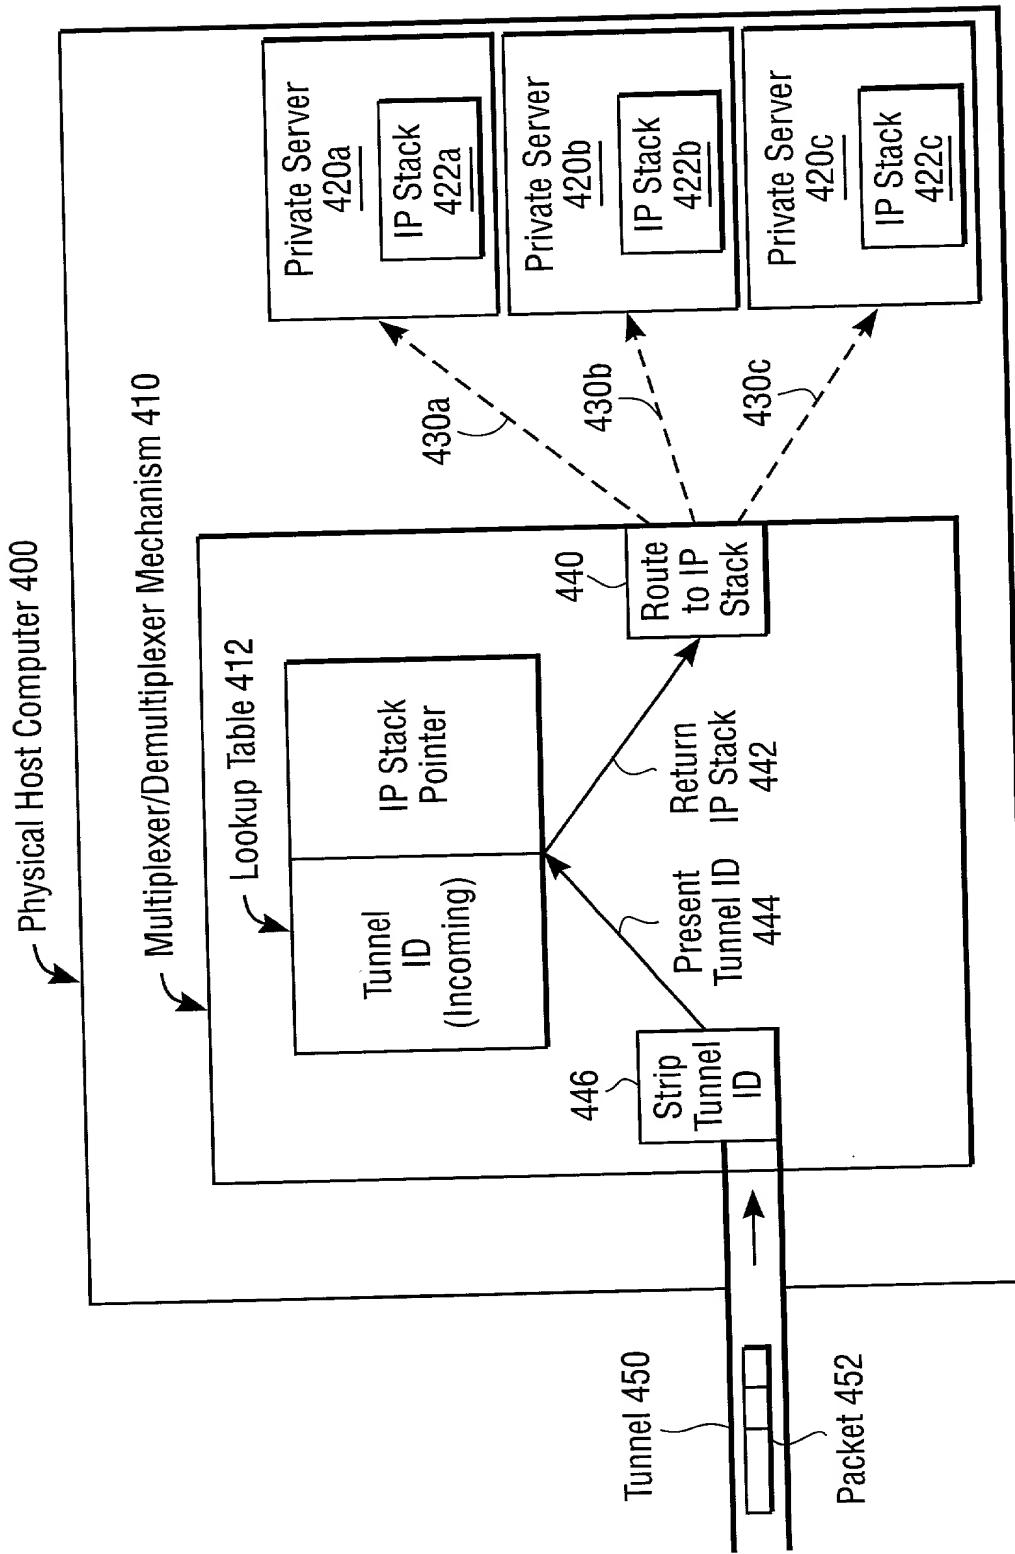
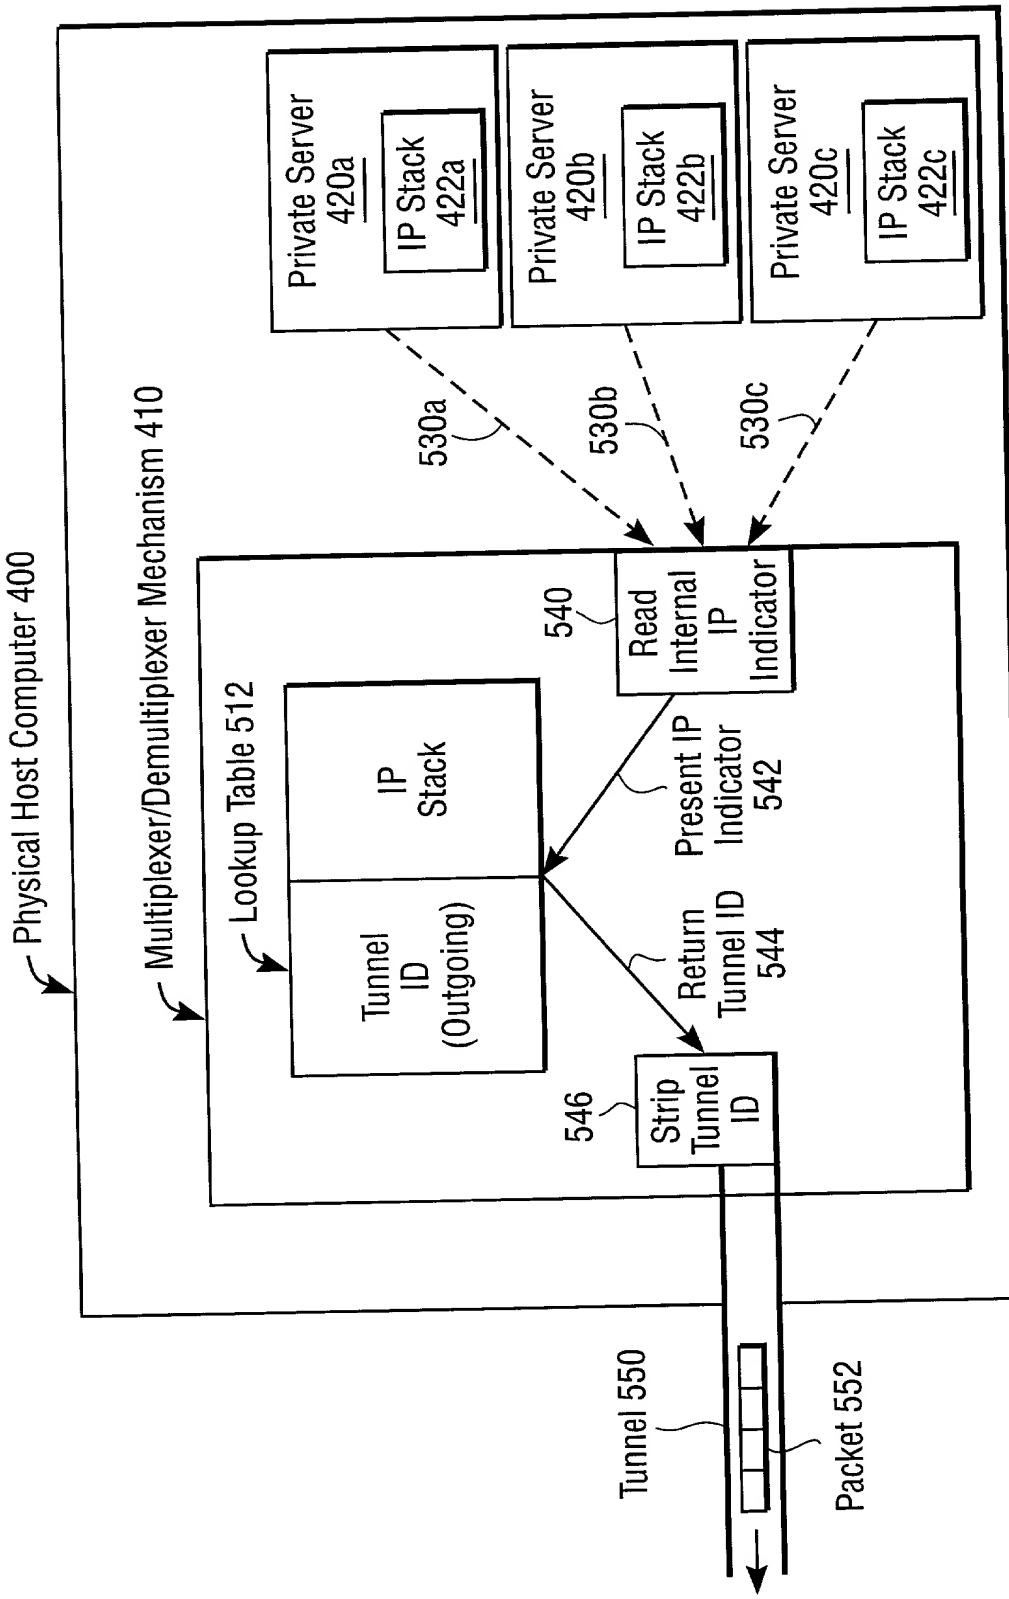


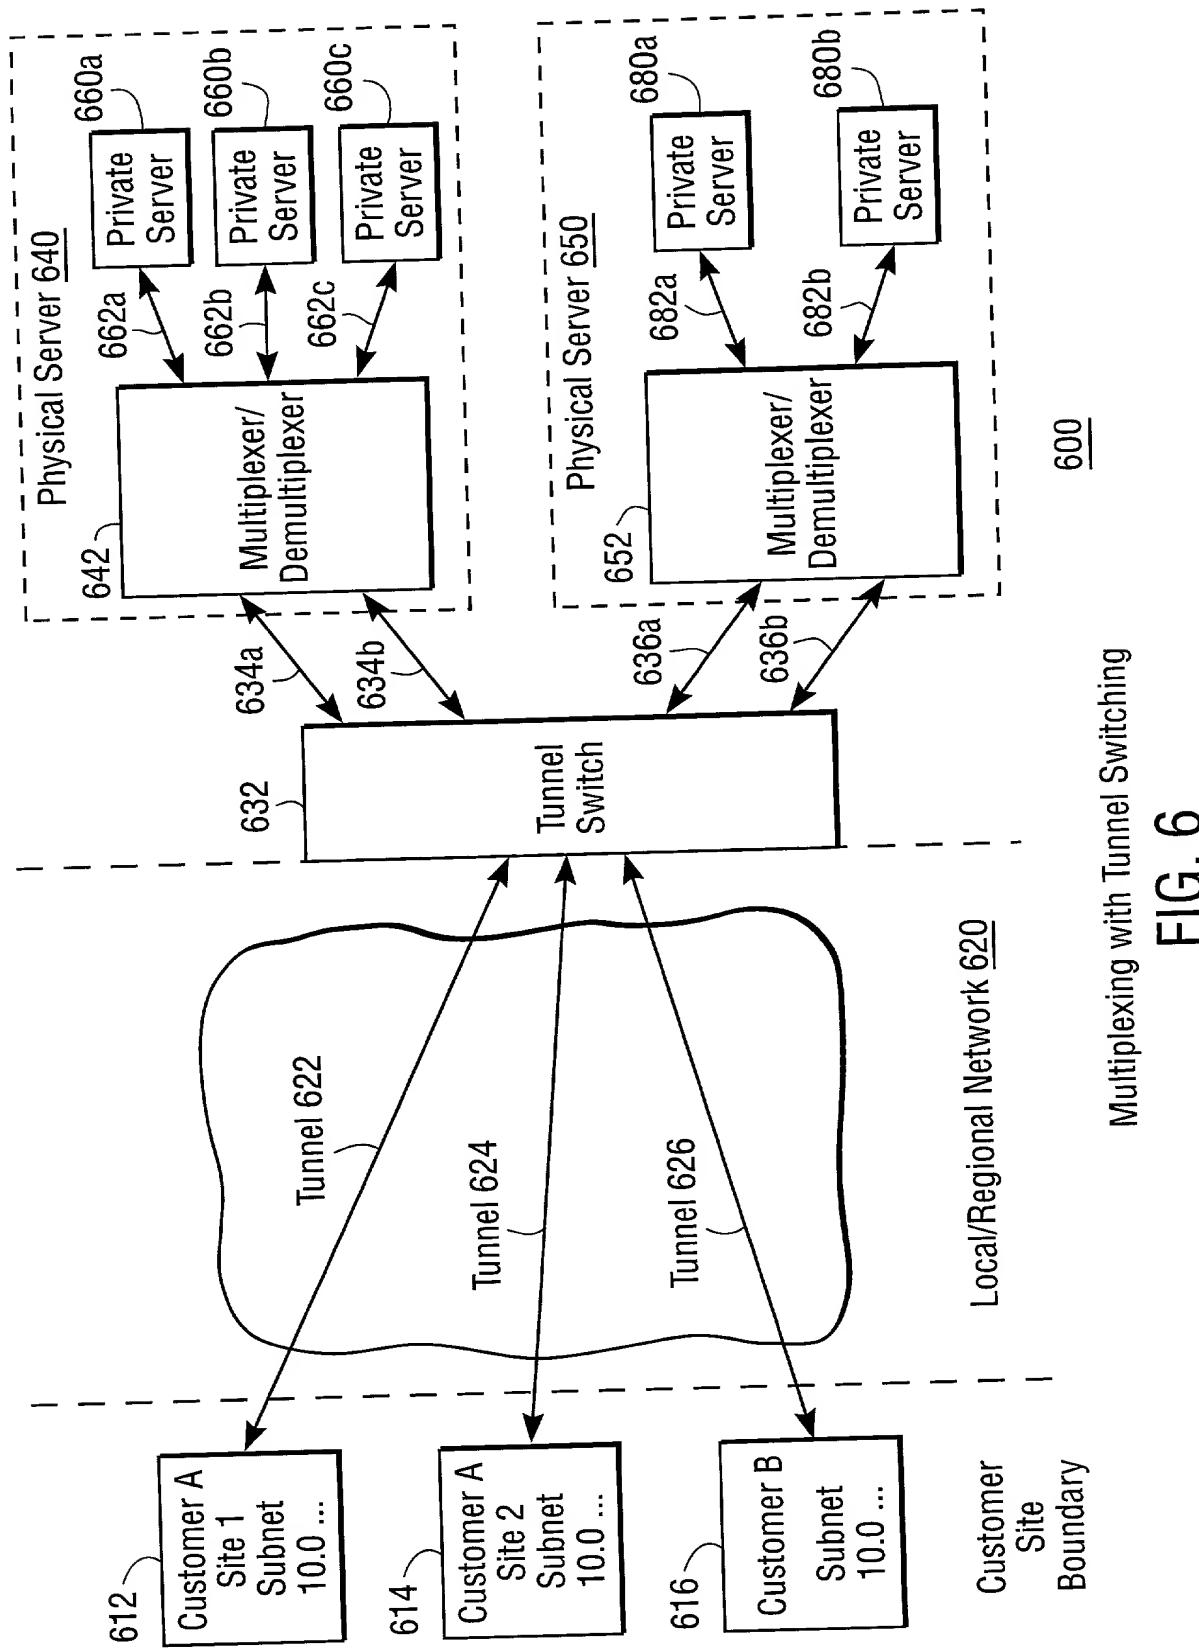
FIG. 4

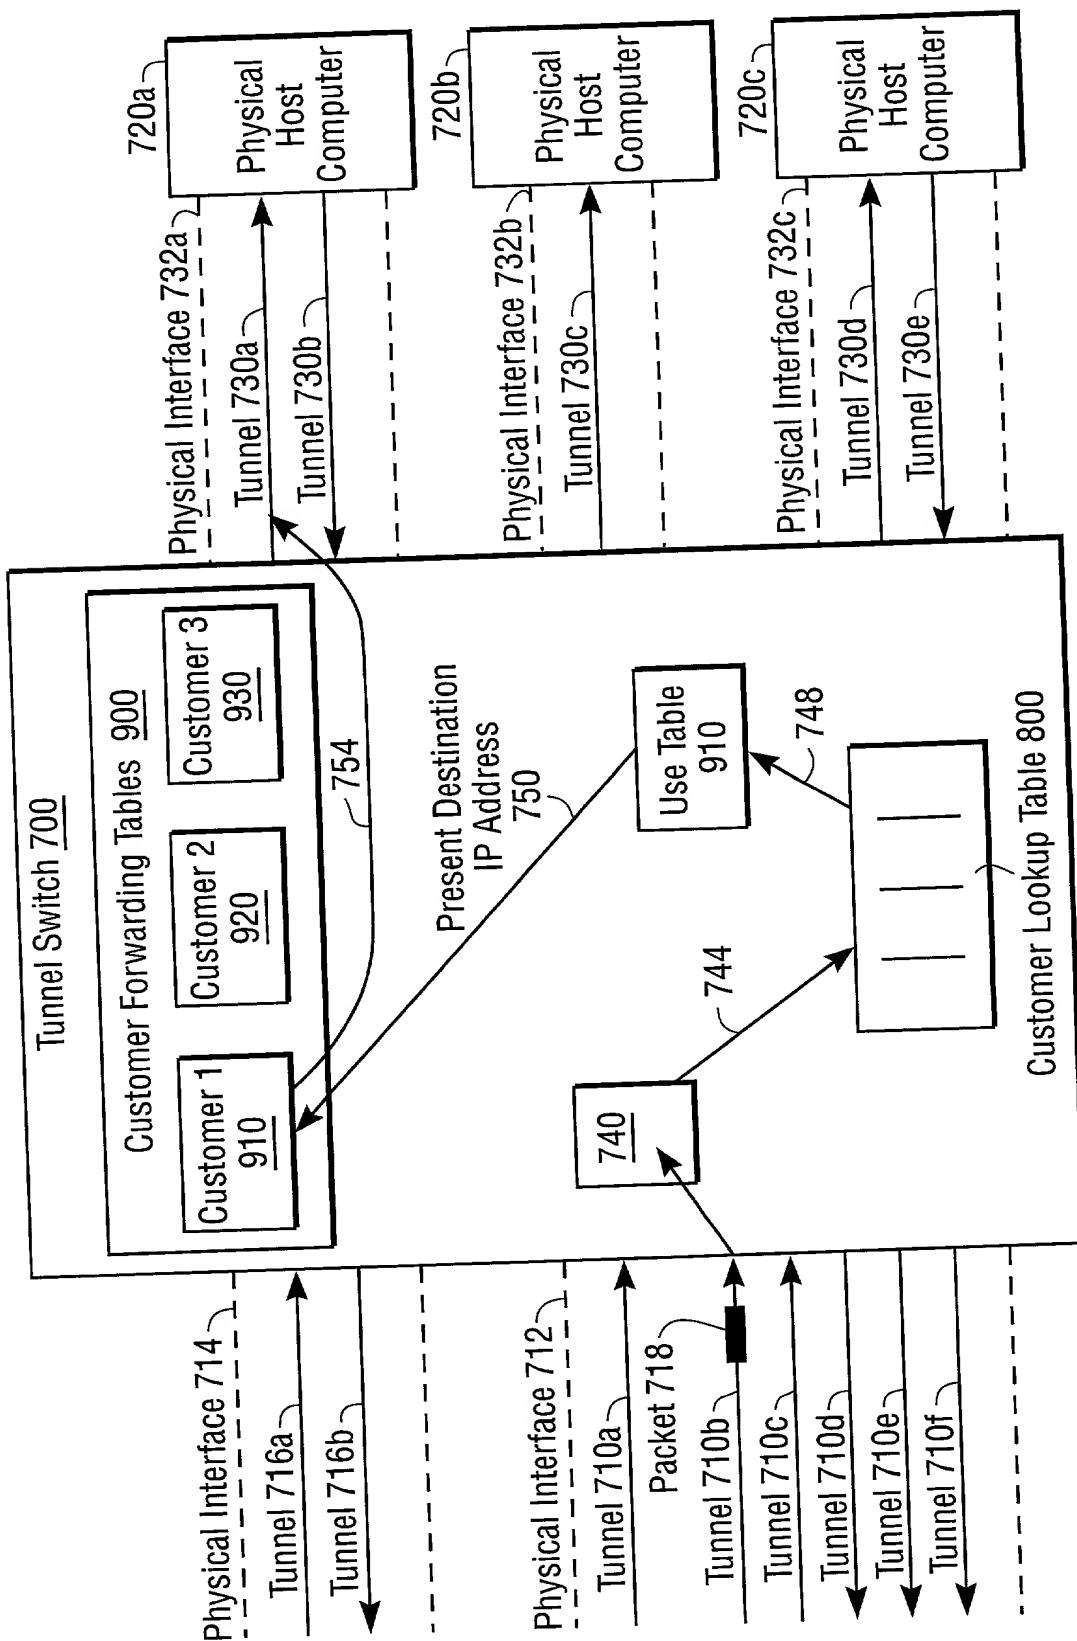
Multiplexer/Demultiplexer
Tunnel-to-Private Server Traffic



Multiplexer/Demultiplexer
Private Server-to-Tunnel Traffic

FIG. 5





Tunnel Switch
FIG. 7

Customer Lookup Table 800

Incoming Physical Interface	Incoming Tunnel ID	Service	Customer ID
712	710a	PS	Customer 1
712	710b	PS	Customer 1
712	710c	TS	Customer 2
732a	730b	PS	Customer 1
732c	730e	PS	Customer 3
		•	•
		•	•
		•	•

Customer Lookup Table

FIG. 8

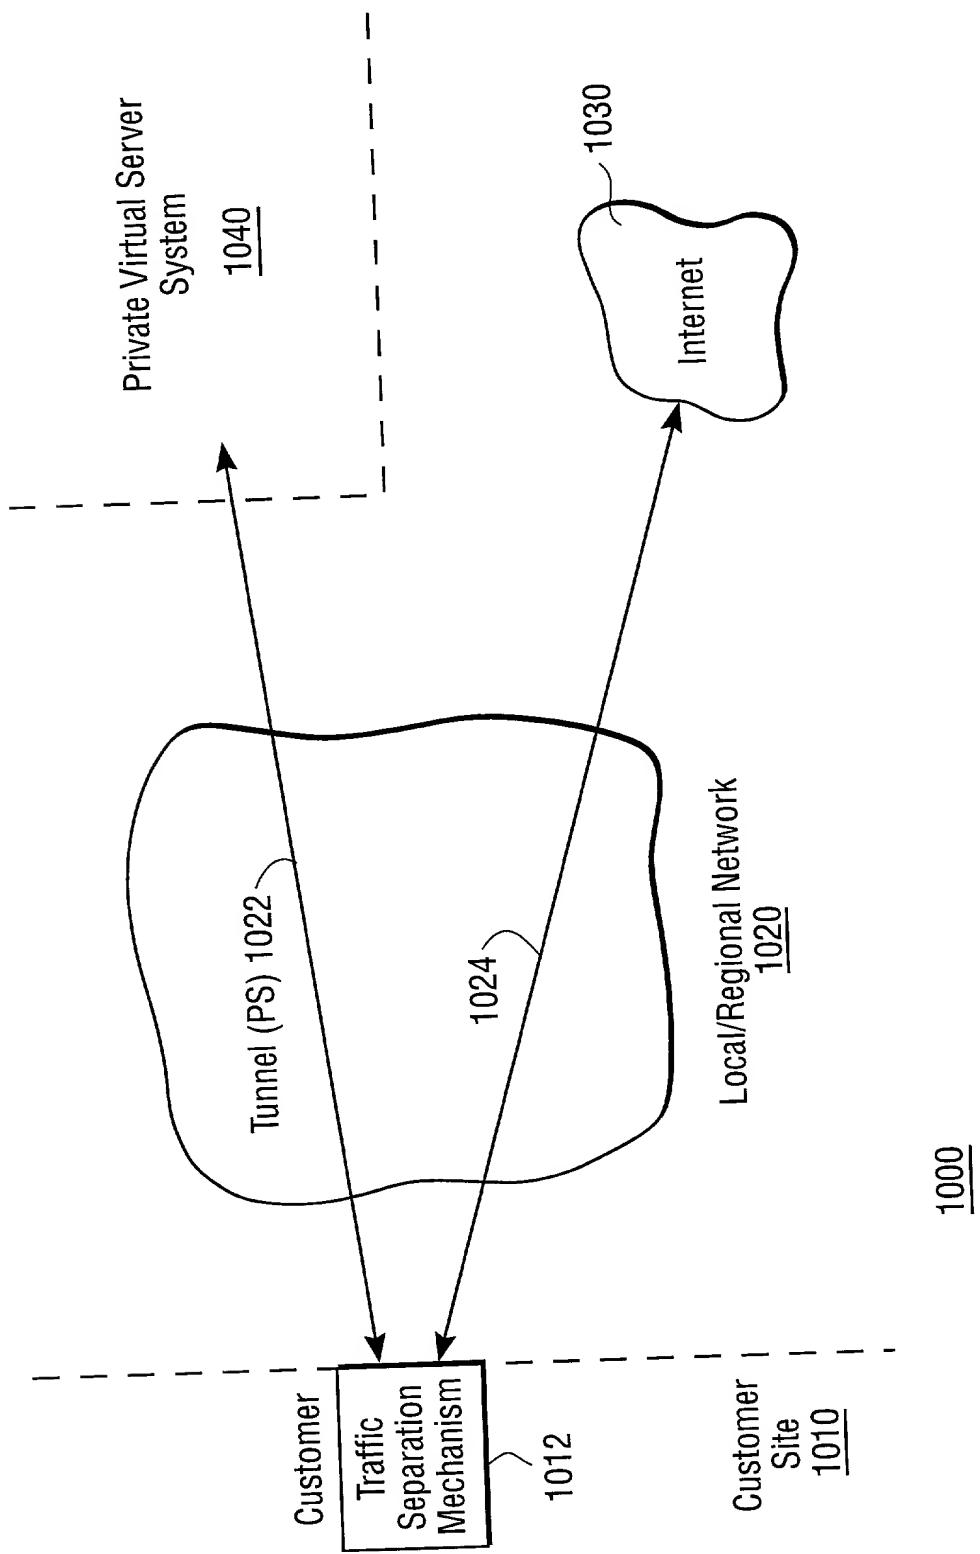
Customer Forwarding Tables 900

Destination IP Address	Outgoing Tunnel ID	Outgoing Physical Interface
Main Server	730a	732a
Backup Server	730c	732b
Site 1	710e	712
Site 2	710f	712
•	•	•
•	•	•
•	•	•

Destination IP Address	Outgoing Tunnel ID	Outgoing Physical Interface
Server	730d	732c
•	•	•
•	•	•
•	•	•

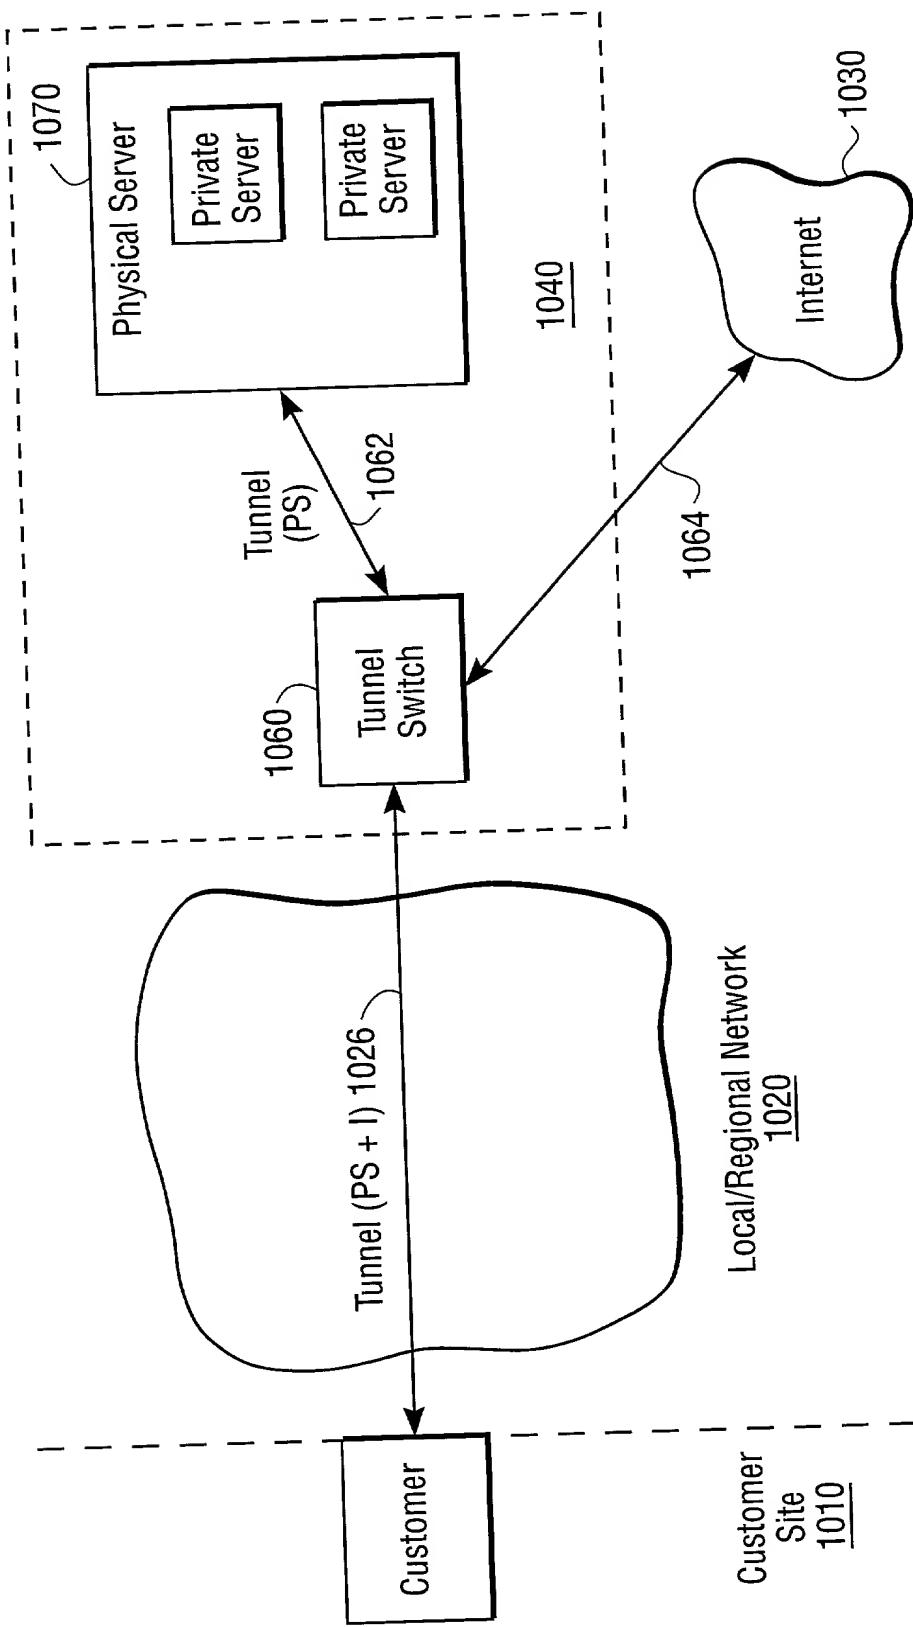
Destination IP Address	Outgoing Tunnel ID	Outgoing Physical Interface
Site 1	710d	712
•	•	•
•	•	•
•	•	•

FIG. 9



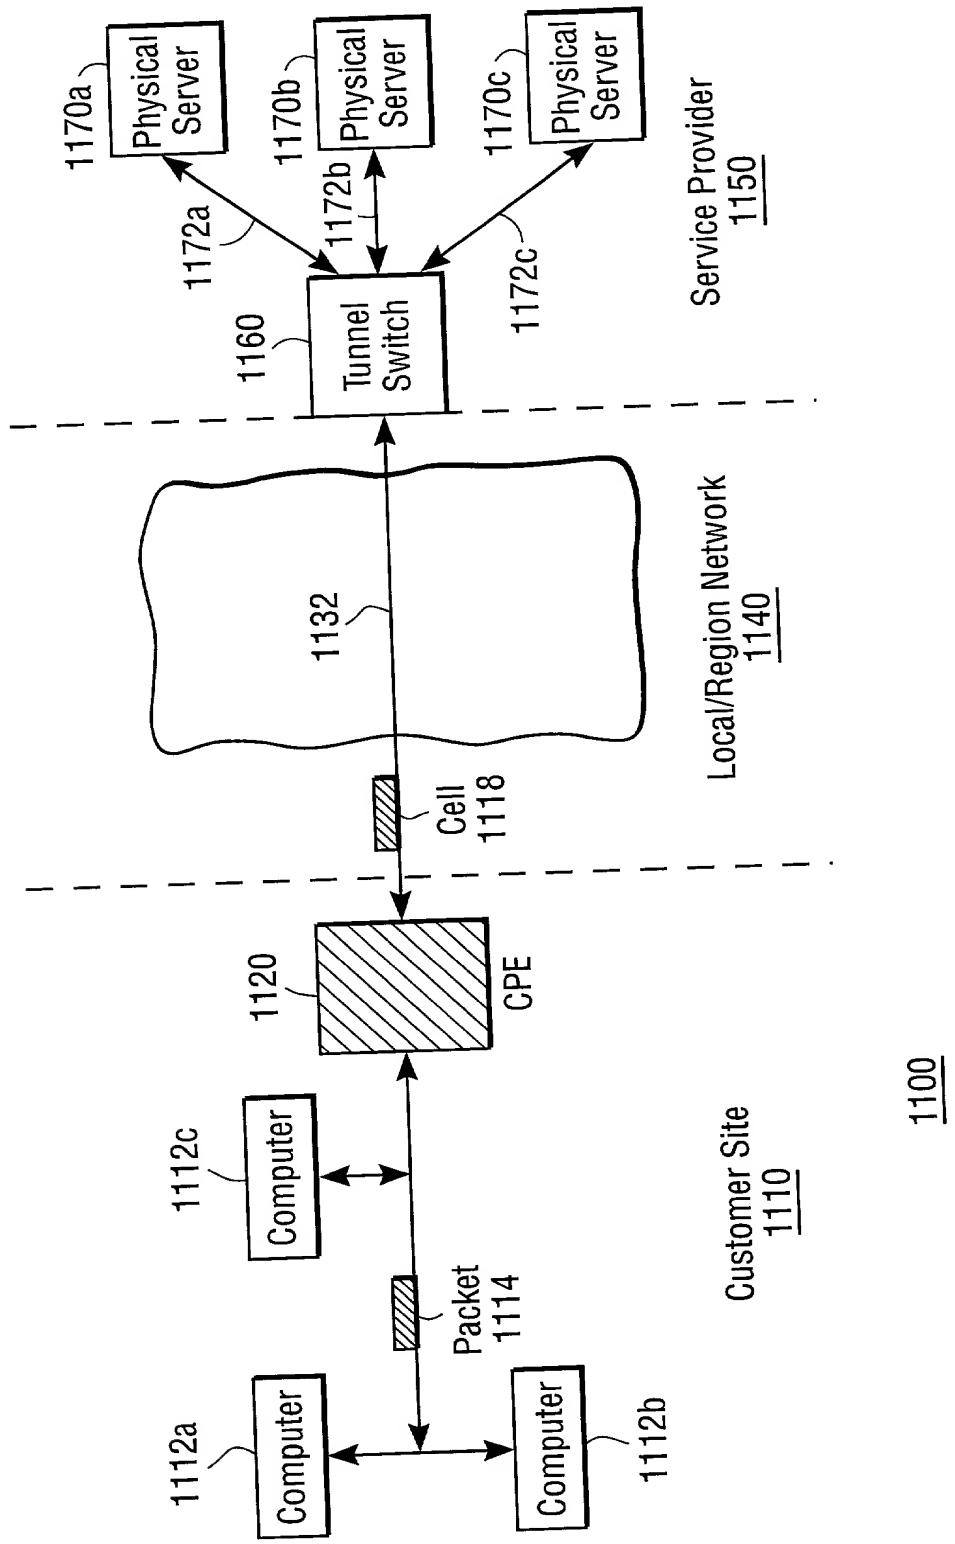
Separation of Internet and Private Server Traffic

FIG. 10A

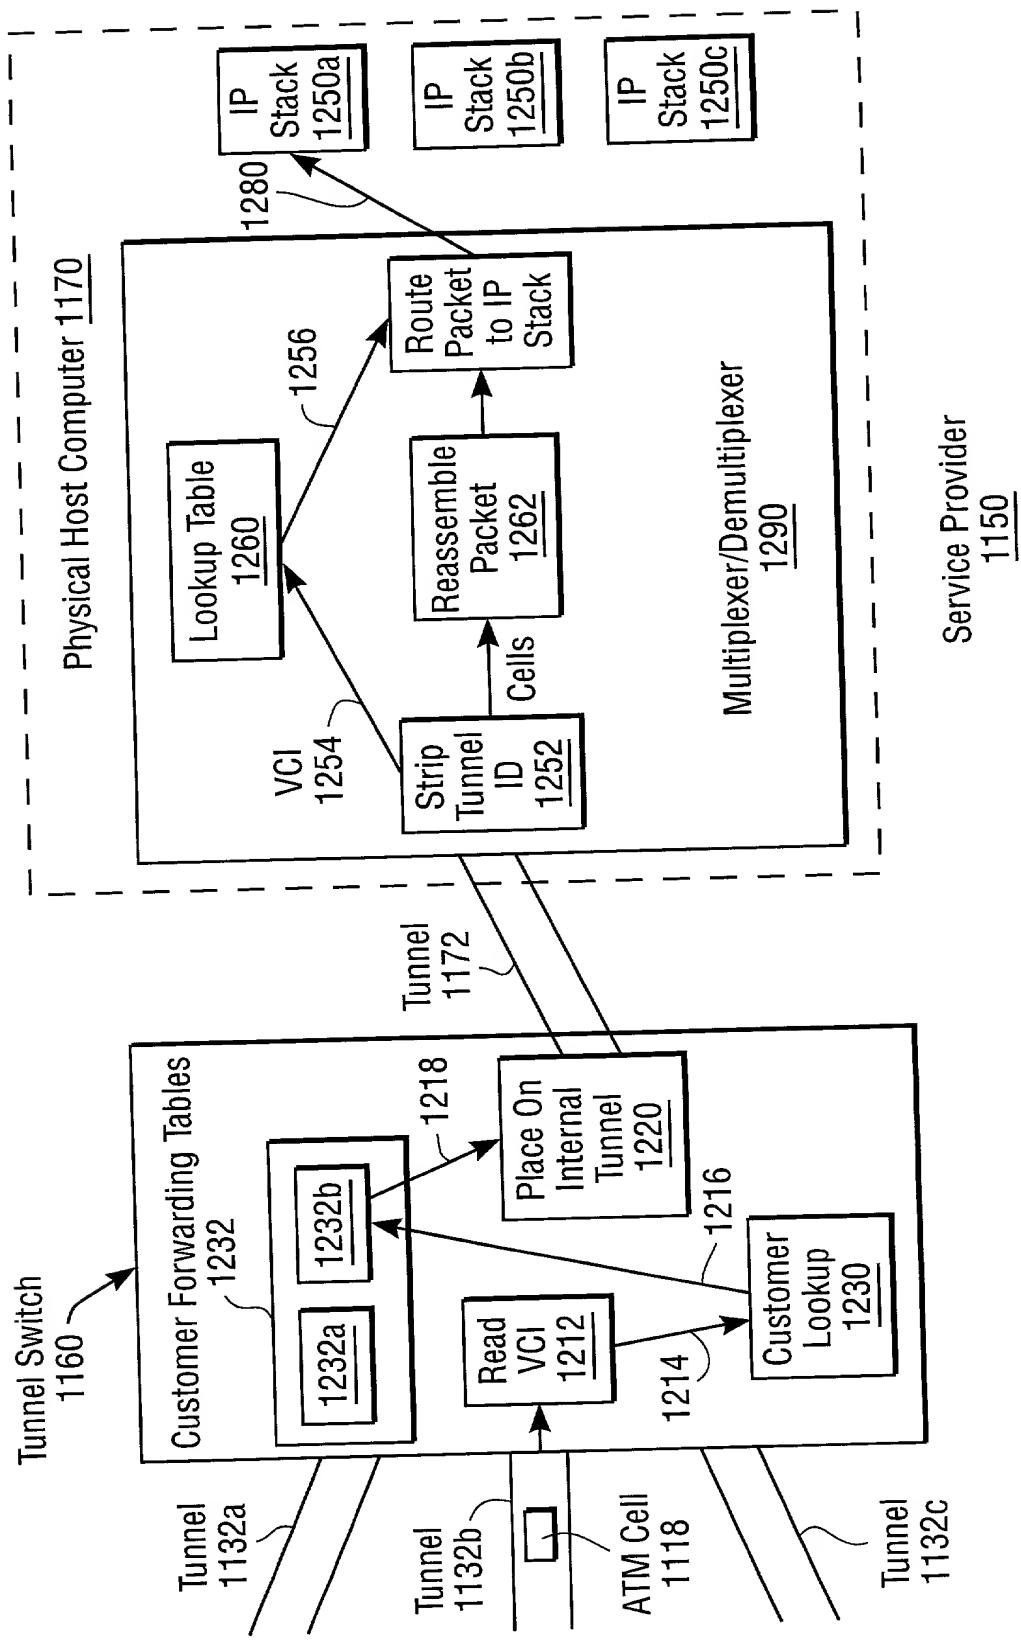


Separation of Internet and PS Traffic

FIG. 10B



ATM Tunnel Using DSL
FIG. 11



ATM Tunnel Trace
FIG. 12

0010/PTO Rev. 6/95	U.S. Department of Commerce Patent and Trademark Office	Attorney Docket Number	4467 US
		First Named Inventor	Peter Newman
<i>COMPLETE IF KNOWN</i>			
		Application Number	
		Filing Date	
		Group Art Unit	
<input checked="" type="checkbox"/> Declaration Submitted with Initial Filing		OR	<input type="checkbox"/> Declaration Submitted after Initial Filing
		Examiner Name	

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES

the specification of which

(Title of the Invention)

is attached hereto

OR

was filed on (MM/DD/YYYY) [] as United States Application Number or PCT International Application Number [] and was amended on (MM/DD/YYYY) [] (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in Title 37 Code of Federal Regulations, § 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code § 119 (a)-(d) or § 385(b) of any foreign application(s) for patent or inventor's certificate, or § 365 (a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? YES	Certified Copy Attached? NO
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Additional foreign application numbers are listed on a supplemental priority sheet attached hereto:

I hereby claim the benefit under Title 35, United States Code § 119(e) of any United States provisional application(s) listed below.

Application Number(s)	Filing Date (MM/DD/YYYY)	<input type="checkbox"/> Additional provisional application numbers are listed on a supplemental sheet attached hereto.

DECLARATION

Page 2

I hereby claim the benefit under Title 35, United States Code § 120 of any United States application(s), or § 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of Title 35, United States Code § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application Number	PCT Parent Number	Parent Filing Date (MM/DD/YYYY)	Parent Patent Number (if applicable)

Additional U.S. or PCT international application numbers are listed on a supplemental priority sheet attached hereto.

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

Name	Registration Number	Name	Registration Number
Robert R. Sachs Laura A. Majerus	42,120 33,417	Albert C. Smith Renee DuBord	20,355 42,500

Additional attorney(s) and/or agent(s) named on a supplemental sheet attached hereto.

Please direct all correspondence to:

Renée M. DuBord
Fenwick & West LLP
Two Palo Alto Square
Palo Alto, CA 94306
U.S.A.

Telephone **(650) 858-7880** Fax **(650) 494-1417**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name of Sole or First Inventor: A petition has been filed for this unsigned inventor

Given Name	Peter	Middle Initial		Family Name	Newman	Suffix e.g. Jr.	
Inventor's Signature	<i>Peter Newman</i>				Date	6 Mar 2000	
Residence: City	Mountain View	State	CA	Country	USA	Citizenship	British
Mailing Address							
Mailing Address	750 N. Shoreline Blvd., #124						
City	Mountain View	State	CA	Zip	94043	Country	USA

Additional inventors are being named on supplemental sheet(s) attached hereto

DECLARATION					ADDITIONAL INVENTOR(S) Supplemental Sheet					
Name of Additional Joint Inventor, if any:			[] A petition has been filed for this unsigned inventor							
Given Name	Pawan		Middle Initial		Family Name	Goyal			Suffix e.g. Jr.	
Inventor's Signature	<i>Pawan Goyal</i>					Date	31/6/00			
Residence: City	Mountain View		State	CA	Country	USA		Citizenship	India	
Mailing Address										
Mailing Address	777 W. Middlefield Road, #83									
City	Mountain View		State	CA	Zip	94043	Country	USA		
Name of Additional Joint Inventor, if any:			[] A petition has been filed for this unsigned inventor							
Given Name			Middle Initial		Family Name				Suffix e.g. Jr.	
Inventor's Signature						Date				
Residence: City			State		Country			Citizenship		
Mailing Address										
Mailing Address										
City			State		Zip		Country			
Name of Additional Joint Inventor, if any:			[] A petition has been filed for this unsigned inventor							
Given Name			Middle Initial		Family Name				Suffix e.g. Jr.	
Inventor's Signature						Date				
Residence: City			State		Country			Citizenship		
Mailing Address										
Mailing Address										
City			State		Zip		Country			
Name of Additional Joint Inventor, if any:			[] A petition has been filed for this unsigned inventor							
Given Name			Middle Initial		Family Name				Suffix e.g. Jr.	
Inventor's Signature						Date				
Residence: City			State		Country			Citizenship		
Mailing Address										
Mailing Address										
City			State		Zip		Country			
[] Additional inventors are being named on supplemental sheet(s) attached hereto										